

AO 120 (Rev. 2/99)

TO: Mail Stop 8 Director of the U.S. Patent & Trademark Office P.O. Box 1450 Alexandria, VA 22313-1450	REPORT ON THE FILING OR DETERMINATION OF AN ACTION REGARDING A PATENT OR TRADEMARK
---	---

In Compliance with 35 § 290 and/or 15 U.S.C. § 1116 you are hereby advised that a court action has been

filed in the U.S. District Court Northern District of California on the following Patents or Trademarks:

DOCKET NO. CV 10-01059 MEJ	DATE FILED 3/12/10	U.S. DISTRICT COURT Northern District of California, San Francisco Division
PLAINTIFF AT&T INTELLECTUAL PROPERTY I, L.P., ET AL.	DEFENDANT TIVO INC.	
PATENT OR TRADEMARK NO.	DATE OF PATENT OR TRADEMARK	HOLDER OF PATENT OR TRADEMARK
1 5,809,492		
2 5,922,045		
3 6,118,976		
4 6,98,478		
5		

In the above—entitled case, the following patent(s) have been included:

DATE INCLUDED	INCLUDED BY <input type="checkbox"/> Amendment <input type="checkbox"/> Answer <input type="checkbox"/> Cross Bill <input type="checkbox"/> Other Pleading	
PATENT OR TRADEMARK NO.	DATE OF PATENT OR TRADEMARK	HOLDER OF PATENT OR TRADEMARK
1		
2		
3		
4		
5		

In the above—entitled case, the following decision has been rendered or judgement issued:

DECISION/JUDGEMENT

CLERK Richard W. Wierking	(BY) DEPUTY CLERK Gloria Acevedo	DATE March 15, 2010
-------------------------------------	--	-------------------------------

Copy 1—Upon initiation of action, mail this copy to Commissioner Copy 3—Upon termination of action, mail this copy to Commissioner
 Copy 2—Upon filing document adding patent(s), mail this copy to Commissioner Copy 4—Case file copy

EXHIBIT A



US005809492A

United States Patent [19]**Murray et al.**

[11] **Patent Number:** **5,809,492**
 [45] **Date of Patent:** **Sep. 15, 1998**

[54] **APPARATUS AND METHOD FOR DEFINING RULES FOR PERSONAL AGENTS**

5,644,686 7/1997 Hekmatpour 395/62

OTHER PUBLICATIONS

[75] Inventors: **LaTondra Alyce Murray**, Raleigh, N.C.; **Loren Gilbert Terveen**, Basking Ridge, N.J.

Pattie Maes, *Agents that Reduce Work and Information Overload*, Jul. 1994, 10 pages.

[73] Assignee: **AT&T Corp.**, Middletown, N.J.

David Canfield Smith, Allen Cypher & Jim Spohrer, *KIDSIM: Programming Agents Without A Programming Language*, Jul. 1994, 13 pages.

[21] Appl. No.: **769,694**

Thomas W. Malone-Kum-Yew Lai and Christopher Fry, *Experiments with Oval: A Radically Tailorable Tool for Cooperative Work*, Nov. 1992, 9 pages.

[22] Filed: **Dec. 18, 1996**

J.C. Sanborn, "A Modifiable Approach to Expert Systems Development," SPIE vol. 786 Applications of Artificial Intelligence V, pp. 99-106, Dec. 1987.

Related U.S. Application Data

[60] Provisional application No. 60/015,070 Mar. 9, 1996.

J.-D. Fekete, et al., "GENESE: Narrowing the gap between Experts and Systems," IEEE Engineering in Medicine & Biology Society 11th Int'l. Conf., vol. 6, pp. 1845-1846, Nov. 1989.

[51] **Int. Cl.⁶** **G06F 15/18**

Primary Examiner—Robert W. Downs

[52] **U.S. Cl.** **706/45; 706/60**

[58] **Field of Search** 395/12, 50, 75, 395/76, 62; 706/45, 60, 11

References Cited**ABSTRACT****U.S. PATENT DOCUMENTS**

4,763,277	8/1988	Ashford et al.	395/62
4,891,766	1/1990	Derr et al.	395/12
5,189,728	2/1993	Yamakawa et al.	395/76
5,283,856	2/1994	Gross et al.	395/50
5,487,135	1/1996	Freeman	395/75
5,566,295	10/1996	Cypher et al.	345/326
5,592,590	1/1997	Jolly	395/75

An apparatus and method for a user to program a personal software agent using an agent manager. The agent manager receives instructions for creating a rule. The rule is placed in a hierarchical order. A determination is made as to whether the rule is valid within the hierarchical order. If the rule is not valid, repairs are suggested to make the rule valid.


30 Claims, 7 Drawing Sheets

FIG. 1




FIG. 2




FIG. 4




FIG. 3




FIG. 5




FIG. 6




FIG. 7




FIG. 8





FIG. 9



APPARATUS AND METHOD FOR DEFINING RULES FOR PERSONAL AGENTS

CROSS-REFERENCES TO RELATED APPLICATIONS

This application claims the benefit of U.S. provisional application Ser. No. 60/015,070, filed Apr. 9, 1996, entitled "Apparatus and Methods for Defining Rules."

BACKGROUND OF THE INVENTION

The invention relates to an agent manager for personal software agents. More particularly, the invention relates to a method and apparatus for defining rules for personal agents to acquire user-specific knowledge and accomplish user 15 specified tasks.

Personal software agents are computer programs that act on behalf of users to perform routine, tedious and time-consuming tasks. To be useful to an individual user, an agent must be personalized to the individual user's goals, habits 20 and preferences. Thus, there exists a substantial need for the agent to efficiently and effectively acquire user-specific knowledge from the user. Previous attempts to solve this problem, however, prove to be unsatisfactory in many ways.

The concept of agency, or the use of agents, is well established. An agent is a person authorized by another person, typically referred to as a principal, to act on behalf of the principal. In this manner, the principal empowers the agent to perform any of the tasks that the principal is unwilling or unable to perform. For example, an insurance agent may handle all insurance matters for a principal, a talent agent may act on behalf of a performer in arranging concert dates, and a foreign agent may spy on other governments to gather secret information on behalf of his or her government.

With the advent of the computer, a new domain for employing agents has opened up. Significant advances in the realm of artificial intelligence have enabled computer programs to act on behalf of computer users to perform routine, tedious and other time-consuming tasks. These computer programs are referred to as "software agents."

Moreover, there has been a recent proliferation of computer and communications networks. These networks permit a user to access vast amounts of information and services without, essentially, any geographical boundaries. Thus a personal software agent has a rich environment to perform a large number of tasks on behalf of a user. For example, it is now possible for a software agent to make an airline reservation, purchase the ticket, and have the ticket delivered directly to a user. Similarly, a software agent could scan the Internet and World Wide Web (WWW) to retrieve information ranging from the latest sports or news to a particular graduate thesis in molecular biology. A scenario could even be envisioned where a personal software agent automatically checks the weather every morning and changes a user's wake-up alarm accordingly, e.g., one-hour later in the case of a snow delay.

To operate effectively, software agents require agent managers to acquire certain information. This information can be generally categorized as: (1) knowledge of the task domains within which the agent is to work; and (2) user-specific knowledge such as a user's preferences, habits and goals.

Knowledge of the task domains within which an agent is to operate is usually supplied by the application developer. Examples of various task domains include electronic messaging, electronic banking or personal financial services.

The fundamental ontology, or vocabulary, usually consists of objects and actions. In messaging, for example, objects includes messages, folders, persons and groups. Actions in messaging include moving a message to a folder, forwarding the message to another user, and replying to the sender. Some objects and actions are very general and thus useful in many domains, e.g., notification actions such as popping up a window on a computer screen, placing a call, and paging.

The agent also requires knowledge about its owner's preferences, habits and goals, which is typically supplied by the user. Much of this type of knowledge can be expressed as If-Then rules. Examples of such rules might include:

1. (messaging) If I get a message from my boss with the subject "urgent" after business hours, then notify me by calling my home number.
2. (banking) If the balance in my checking account falls below \$500 and the balance in my savings account is above \$2000, then transfer \$500 from savings to checking.
3. (finances) If the price of XYZ stock falls below \$30, buy 100 shares.

The effective acquisition of user-specific knowledge has posed a serious problem in the area of personal software agents. Conventional methods for a software agent managers to acquire user-specific knowledge fall within three broad categories: (1) learning; (2) programming by demonstration (PBD), and (3) end-user programming (EUP). Each of these categories of methods will be discussed in turn.

The learning approach is detailed in a paper by P. Maes titled "Agents that Reduce Work and Information Overload," *CACM* 37, 7 (July 1994), pp. 55-67. The agent learns rules by watching user behavior and detecting patterns and regularities. The agent records a user's activity in terms of situation-action pairs. For example, a situation 25 could be an e-mail message and an action the way the user processed the message. The system predicts how a new situation should be handled by finding the most similar previously seen situation. After seeing enough examples, the system is able to make good predictions, and make suggestions to the user. The user may also instruct the agent to act 30 on its own when the user's confidence in the agent's predictions exceeds a specified threshold.

The problem with the learning method is that the agent manager has a slow learning curve. The agent manager is 45 unable to offer any advice until it has seen a significant number of examples. Further, the agent manager cannot offer assistance in truly new situations since there were no previous examples from which to draw. This limitation can be minimized by permitting one software agent manager to 50 access the experience of another software agent manager. This solution, however, fails to personalize the agent in the preferences, habits and goals of the agent's owner.

The PBD method is discussed in a book by A. Cypher titled "Watch What I Do: Programming by Demonstration," 55 MIT Press, Cambridge, Mass., 1993. In PBD, a user puts a system into record mode to record a user's actions in an executable file. This technique is similar to the macro function on many popular word processors. The prototypical PBD systems operate in domains that have a natural graphical representation. One example is KIDSIM which is outlined in a paper by D.C. Smith et al. titled "KIDSIM: Programming Agents Without a Programming Language," *CACM* 37, 7 (July 1994), 55-67. In KIDSIM, children create simulations in which characters move around in a two-dimensional world. They create rules by using a graphical editor to depict changes of state—that is, they draw the configuration of the world before and after the rule has 60 65

3

applied. The system generalizes from the concrete objects in the demonstration to create an abstract rule, e.g., replacing a particular character by any character of a certain type and replacing absolute locations by relative changes in location.

The problem with the personal agents using the PBD method is that there is no mechanism by which the personal agent may determine if rules are incompatible with other rules. Moreover, if there are incompatibilities, the personal agent does not offer solutions for repairing the invalid rule.

The foundational research on end-user programming (EUP) is found in a work by T. W. Malone et al. titled "Experiments with Oval: A radically Tailorable Tool for Cooperative Work," *ACM Transactions on Information Systems* 13, 2 (April 1995), 177-205. A user first indicates the type of object to which the rule applies, such as a message. The user then describes the conditions of the rule by filling out a form that contains a field for each attribute of the object. For messages, the attributes would include sender, subject, recipient, and so forth. Finally, the user specifies the actions of the rule by selecting from a list of all the actions that could apply to this type of object. For example, a message could be moved, copied, forwarded, deleted, and so on.

As with personal agents using the PBD method, traditional personal agents using the EUP method are incapable of determining whether rules are incompatible with other rules, and if so, offering repairs to the rule.

A substantial need, therefore, exists for an agent manager which can acquire user-specific knowledge in the form of rules in a manner which is timely and precise, and which can detect and repair any conflicts with existing rules.

SUMMARY OF THE INVENTION

In view of the foregoing, there exists a need in the art for a method and apparatus for solving the above-stated problems. More particularly, there is a need in the art for a method and apparatus for a personal software agent using enhanced EUP methods to acquire knowledge about a user's habits, preferences and goals using an agent manager. The agent manager allows the user to define new rules, and determines whether such rules conflict with existing rules. If such a conflict exists, the agent manager interactively guides the user in selecting and applying repairs for the new rules to remove the conflict.

Thus, one embodiment of the invention permits a user and agent manager to collaborate in developing and managing a set of rules that embody the user's preferences for handling a wide variety of situations.

This is accomplished using an apparatus and method for defining rules for a personal software agent. One embodiment of the invention includes an agent manager which uses a rule edit module for defining at least one condition and at least one action to form a new rule. The invention uses a rule index module for placing the new rule in a rule hierarchy, wherein the hierarchy is comprised of parent rules, child rules, and sibling rules. The rule index module adds actions from the parent rules to the new rule and from the new rule to the child rules. The rule index module also creates intersection rules from an intersection of the new rule with the sibling rules. The invention uses a rule analysis module for determining whether the new rule and the intersection rules conflict with any existing rules in the hierarchy. The invention uses a rule repair module for repairing the new and intersection rules where such conflicts exist.

With these and other advantages and features of the invention that will become hereinafter apparent, the nature

4

of the invention may be more clearly understood by reference to the following detailed description of the invention, the appended claims and to the several drawings attached herein.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a flow diagram illustrating the overall structure of one embodiment of the invention.

FIG. 2 is a flow diagram of a decision-making portion of a rule edit module used in accordance with one embodiment of the invention.

FIG. 3 is a picture of a rule editor used in accordance with one embodiment of the invention.

FIG. 4 is a flow diagram of a new rule portion of a rule edit module used in accordance with one embodiment of the invention.

FIG. 5 is a flow diagram of an existing rule portion of a rule edit module used in accordance with one embodiment of the invention.

FIG. 6 is a flow diagram of a rule index module used in accordance with one embodiment of the invention.

FIG. 7 is a flow diagram of a rule analysis module used in accordance with one embodiment of the invention.

FIG. 8 is a flow diagram of a first portion of a rule repair module used in accordance with one embodiment of the invention.

FIG. 9 is a flow diagram of a second portion of a rule repair module used in accordance with one embodiment of the invention.

DETAILED DESCRIPTION

Referring now in detail to the drawings wherein like parts are designated by like reference numerals throughout, there is illustrated in FIG. 1 a flow diagram illustrating the overall structure of one embodiment of the invention. FIG. 1 shows four modules: rule edit module 200; rule index module 300; rule analysis module 400, and rule repair module 500.

The invention relates to an agent manager for personal software agents. More particularly, the invention relates to a method and apparatus for defining rules for personal agents to acquire user-specific knowledge and accomplish user specified tasks.

The invention comprises an agent manager which improves end-user programming by allowing the agent manager to perform the more complicated programming tasks for the user, thereby minimizing programming commands and syntax for the user. The agent manager allows the user to focus on three tasks: (1) define conditions; (2) define attributes; and (3) define actions. Once these three tasks are accomplished, a rule is formed. As a user creates or edits rules, the agent manager analyzes the rules for problems, computes methods of repairing the problems, and guides the user in selecting and applying the repairs.

The agent manager's role is to watch for problems as users create or edit rules and help users to repair problems that it detects. Specifically, the agent manager detects when a rule's plan is not valid, i.e., some action does not apply or the rule conflicts with an existing rule. This can occur either as a user edits a rule or when a new rule is added to a rule hierarchy (due to interactions with other rules, described more fully below). When the agent manager finds an invalid plan, it computes repairs, which are changes to the plan that will make the plan valid. It then presents an interactive explanation that explains the nature of the problem and

guides the user in selecting a repair. One embodiment of invention will be discussed in detail using electronic mail as an example.

FIG. 2 is a flow diagram of a portion of a rule edit module used in accordance with an embodiment of the invention. The invention uses rule edit module 200 to specify rule conditions and actions. As shown in FIG. 2, if a user wants to enter a new rule at step 202, the rule edit module performs the actions detailed in FIG. 4. If a user wants to edit an existing rule at step 202, the rule edit module performs the actions detailed in FIG. 5.

FIG. 3 is a picture of a rule editor used in accordance with one embodiment of the invention. The following descriptions of FIG. 4 and FIG. 5 may be enhanced by referring to FIG. 3.

FIG. 4 is a flow diagram of a new rule portion of a rule edit module used in accordance with one embodiment of the invention. At step 204, a user specifies rule conditions, examples of which include restricting the values of the message type, sender, recipient, subject, and the date and time of arrival. By specifying these rule conditions, the user defines the set of messages to which the rule applies. For example, FIG. 3 shows a rule that applies to e-mail messages from Julia Hirschberg whose subject field contains the word "important." As the user specifies the conditions, the system continuously updates and displays the rule's context, i.e., its relationship to other rules and actions the rule will inherit from more general rules (as discussed in reference to FIG. 6).

The user can also specify using rule edit module 200 two types of conditions about the arrival date and time of messages. First, users may specify messages that arrive in a particular interval, such as "between 4:00 p.m. on Sep. 15, 1995 and 8:00 a.m. on Sep. 20, 1995." The user can also specify intervals such as "from November 1995 to March 1996," and "after 5:00 p.m. on Sep. 15, 1995." Second, the user can specify messages that arrive during a particular period, such as "during June, July and August," "on Saturday or Sunday," "between 1800 and 0800," "on the first Monday of each month," and "on the last working day of the month."

As the user enters the conditions at step 204, these conditions define an initial state. The initial state is the class of messages to which the actions apply. For example, the initial state for the rule in FIG. 3 is

InSpool and
messageType = email and
sender = Julia-Hirschberg and
subject contains "important".

This initial state will be transformed as new actions add or delete facts from the initial state.

At step 208, a user selects rule actions. In our e-mail example, the set of actions could include Move Message, Forward Message, Notify by Paging, Notify with Pop Alert, and so forth. As shown in FIG. 3, the rule contains a single action that moves messages to a folder named "julia." The system has determined that one rule named "Email" is more general, and that one action "Copy Message 'email archive, '" will be inherited to this rule (as discussed in reference to FIG. 6).

At step 210, the agent manager tests whether the selected action is permitted by the system in the state in which the specified conditions and previous actions have been applied. If it is, the agent manager updates the current state at step

220. If the action is not permitted, the user is informed and can then select another action, or change the rule conditions.

FIG. 5 is a flow diagram of an existing rule portion of a rule edit module used in accordance with one embodiment of the invention. As shown in FIG. 2, if the user wants to edit an existing rule, the agent manager performs the actions detailed in FIG. 5.

As shown in FIG. 5, the agent manager tests whether there is any change in rule conditions at step 222. If there is, the agent manager redefines the initial state at step 224 and must recheck the validity of the entire plan (as detailed in reference to FIG. 7). The agent manager sets a condition flag at step 226, and tests whether the user has changed any actions at step 228. If at step 222 no changes in the conditions have been detected, the agent manager moves directly to test whether the user has changed any actions at step 228. If the user has not changed any actions at step 228, the agent manager checks whether the condition flag has been set at step 248. If it has, the agent manager resets the condition flag at step 238 and moves to rule analysis module 400. If the condition flag has not been set, the program waits for user input, or the termination of the agent manager.

If the user changes an action at step 228, the agent manager tests whether the user added an action at step 230. If the user has added an action to the end of plan, all that must be checked is that action's applicability, i.e., the action is possible for the selected condition. Thus, at step 232 the added action is tested as to whether it is valid. If it is not, the user is prompted or directed to select another action. Alternatively, the user could select a condition that conforms to the action initially selected. If the action is valid, the agent manager updates the current state of the plan at step 234 and checks to see if the user changed any conditions by checking whether the condition flag was set at step 236. If no conditions were changed, this program module ends, or alternatively, allows the user to make additional edits to the existing or other rule plans. If the condition flag is set at step 236, then the condition flag is reset at step 238 and the program module returns.

If at step 230, the user did not add an action, the agent manager tests whether the user wants to delete or insert an action at steps 240 and 242, respectively. If either action is selected, the agent manager checks at step 244 whether the action is valid. If it is, the agent manager updates the current state of the plan at step 246, resets the condition flag at step 238, and returns. If the action is not valid at step 240 or 242, the program informs the user and prompts the user to perform another action.

FIG. 6 is a flow diagram of a rule index module used in accordance with one embodiment of the invention. To better understand how the agent manager indexes rules, it may be helpful to first describe the representation and reasoning technology used by the invention.

The rule representation used in the invention should enable three functions:

1. Determine whether an action may apply in a given context, e.g., to a specified set of messages, or after certain other actions;
2. Determine when one rule subsumes another, that is determine whether the set of objects to which the first rule applies is a superset of the set of objects to which the second rule applies;
3. Determine when two rules intersect, that is determine whether the set of objects to which the first rule applies intersects the set of objects to which the second rule applies.

One example of representing rules in accordance with the invention is through the use of CLASSIC, which is detailed

in a paper by A. Borgida et al. titled "CLASSIC: A Structural Data Model for Objects," *SIGMOD '89*, (1989), and a paper by R. J. Brachman titled "Living with Classic: When and How to Use a KL-ONE-Like Language," *Formal Aspects of Semantic Networks*, J. Sowa, Ed., Morgan Kaufmann, Los Altos, Calif., 1990, both of which are herein incorporated by reference in their entirety.

CLASSIC is a description logic. As such, it permits users to create structured descriptions of sets of objects (known as concepts) and individual objects. As with any standard knowledge representation or object-oriented language, users may state hierarchical relationships between concepts. The main service provided by CLASSIC, however, is to determine subsumption relationships automatically by analyzing concept descriptions.

Determining subsumption relationships is a key requirement for the invention. The invention uses CLASSIC to automatically maintain a rule hierarchy. For example, the conditions of the rule being edited in FIG. 3 would be represented in CLASSIC as the following concept (referred to as C_0):

C_0 - messageType = email and
sender = Julia-Hirschberg and
subject contains "important"

Examples of concepts that subsume C_0 include concepts formed from a subset of its conjuncts, such as "messageType=email" (the more general rule shown in FIG. 3) or "sender=Julia-Hirschberg." The object hierarchy is used, so "messageType=Message" also subsumes C_0 . In addition, the semantics of substrings are used, so a concept like "subject contains 'port'" also subsumes C_0 .

CLASSIC also determines the intersection between two concepts. This is useful because users often state general rules that categorize messages by a single attribute, such as:

sender = Julia-Hirschberg
subject contains "important."

CLASSIC determines whether these rules intersect. The intersection between these two rules is the conjunction of the two concepts. In general, however, CLASSIC may have to combine parts of two concepts to determine the intersection. For example, consider the two descriptions of message arrival times "between 12:00 on Nov. 8, 1995 and 12:00 on Nov. 10, 1995" and "between 18:00 and 8:00." CLASSIC determines that the intersection is "18:00 to 23:59 on Nov. 8, 1995, midnight to 8:00 and 18:00 to 23:59 on Nov. 9, 1995, and midnight to 8:00 on Nov. 10, 1995."

Like rule conditions, rule actions are represented as CLASSIC concepts. Actions are described in terms of a basic artificial intelligence (AI) planning model. Each action applies in a particular "state of the world," represented as a list of facts, and each action may change the state of the world. Thus, an action is defined in terms of its preconditions (facts which must be true for the action to be applied), add list (facts added by the actions), and delete list (facts deleted by the action). For example, the action PickupCall is defined as:

preconditions: TelephoneCall
add: CallPickedup

If an action A_1 occurs in a plan before an action A_2 , and A_1 deletes a precondition of A_2 , then A_1 defeats A_2 .

AI planning research explores planning synthesis (creating a plan that transforms a specified initial state to a specified final state) and recognition (inferring an agent's plan by observing his/her/its actions), both of which are very difficult computational problems. By way of contrast, the invention allows the user to construct a plan, i.e., the actions of a rule. Thus the agent manager need only check plan validity, which is a much less expensive (in terms of CPU cycles) computational task.

Thus, as shown in FIG. 6, the agent manager inserts the new rule defined by the user into a rule hierarchy at step 302. One way to establish a rule hierarchy is to use CLASSIC. The plan from parent rule PARENT is inherited to the new plan NEW and added at the proper position in NEW's plan at step 304. NEW's plan is inherited to its child CHILD and inserted in the proper position, i.e., after the local plan of CHILD and before the more general plan of PARENT, at step 306. The intersection INTER of NEW and its sibling rule SIBLING is formed at step 308, and all the appropriate plans are inherited to INTER at step 310. It is worthy to note that the user only becomes aware of INTER if some action in its plan is defeated. All the composed plans are checked for validity by rule analysis module 400, and whenever any actions of rule's plan are defeated, the display of the rule is updated to reflect this fact.

FIG. 7 is a flow diagram of a rule analysis module used in accordance with one embodiment of the invention. The rule analysis module 400 of the agent manager detects when a rule's plan is not valid, i.e., some action does not apply. This can occur either as a user edits a rule or when the rule is added to the hierarchy. When the agent manager finds an invalid plan, it computes repairs, which are changes to the plan that will make the plan valid. It then presents an interactive explanation that explains the nature of the problem and guides the user in selecting a repair.

The task of verifying plan validity centers around the notion of current state. The conditions of a rule define an initial state that is the class of messages to which the actions apply. Each action transforms the initial state by adding or deleting facts. For example, the initial state for the rule in FIG. 3 is:

InSpool and
messageType = email and
sender = Julia-Hirschberg and
subject contains "important."

InSpool is a special tag indicating that the message has not yet been disposed of, i.e., deleted or moved to a folder. All actions have at least one precondition—InSpool—which is the one fact in delete list of "Move Message." Thus, adding a "Move Message" action results in the removal of InSpool from the current state. If any other actions are subsequently added to the rule, the system will find that they do not apply.

The system has to update the current state whenever an action is added, deleted or inserted, or when the conditions are changed, as described previously in reference to FIG. 5. When an action is added to the end of the plan, all that has to be checked is that action's applicability. Changing the conditions redefines the initial state, so the validity of the whole plan must be rechecked. When an action is inserted into or deleted from the middle of the plan, the whole plan again is checked. Alternatively, the checking could begin from the insertion/deletion point by keeping a record of previous states.

As a user edits a rule, the agent manager continuously recomputes its position in the rule hierarchy, thus determin-

ing the actions to inherit to the rule from more general rules. Either the locally defined or the inherited actions of a rule may be done first. If the user has not specified an order, the system will prefer the local-first ordering, but if this results in an invalid plan, it will try inherited-first. For example, in FIG. 3, if the local "Move Message" action were ordered before the inherited "Copy Message" action, the plan would be invalid, so the invention selected the inherited-first ordering.

When a rule is added to the hierarchy, the Agent Manager inherits its plan to all more specific rules. It also considers rules resulting from the intersection of the new rule and its sibling rules. These cases all result in new plans being composed, all of whose validity must be tested. Notice that users become aware of intersection rules only if their plans are invalid.

After actions are inherited, the plan at each rule to which actions have been inherited must be checked for validity. These rules are (potentially) NEW, all descendants of NEW, and all INTER rules of NEW and its siblings.

Thus, as shown in FIG. 7, rule analysis module 400 compares NEW, CHILD and INTER conditions with conditions from existing rules OLD at step 402. If a NEW, CHILD or INTER condition and OLD condition are the same, or simultaneous, at step 404, the agent manager determines if the actions of NEW, CHILD or INTER and OLD conflict (as previously discussed). If they do conflict, the NEW or INTER plan is invalid, and rule analysis module 400 moves to the next plan at step 410. If all plans are analyzed, risk analysis module 400 returns. If not, rule analysis module 400 processes the next plan from NEW, CHILD or INTER.

At step 404, if the conditions of NEW, CHILD or INTER and OLD are not the same, the agent manager tests whether the conditions of NEW, CHILD or INTER and OLD intersect at step 414. If they do, the agent manager checks if the actions conflict at step 406. If the actions conflict, the NEW, CHILD or INTER plan is invalid at step 408, and rule analysis module 400 moves to the next plan. If at step 414 the conditions of NEW, CHILD or INTER and OLD do not intersect, the plan is considered valid at step 412.

FIG. 8 is a flow diagram of a first portion of a rule repair module used in accordance with one embodiment of the invention. When the agent manager detects that a rule has an invalid plan at step 502, it attempts to categorize the rule as one of a pre-defined set of invalid rule types at step 506. For example, the intersection of a rule that files messages by sender (e.g., "move messages from Tom to the folder Tom-Mail") and a rule that files messages by subject (e.g., "move messages with the subject 'RDDI' to the folder RDDI-Project") constitutes a common type of problem. The benefit of categorizing problems is that multiple instances of a given category may occur as a user works, and the user may specify that all problems of a particular category should be repaired in the same way.

When an action in a plan does not apply, the agent manager computes repairs at step 508, and generates an interactive dialogue that explains why the action does not apply and suggests repairs to the plan that will enable the action at step 510. How the agent manager computes repairs for a rule at step 508, and creates the interactive dialogue at step 510, is described in more detail as follows.

The dialogue contains a component for each unsatisfied precondition of the action. The text is created by composing templates associated with actions, preconditions and plan modification operators. If a precondition was deleted by a previous action in the plan, the system looks for alternative actions whose add list is a superset of the defeating action and whose delete list does not contain the unsatisfied precondition. Consider the rule formed from the intersection of:

5

messageType = email and
sender = Julia-Hirschberg
=> Move Message "julia", and
messageType = email and
subject contains "important"
=> Move Message "important".

10

The two "Move Message" actions conflict with each other, since each deletes "InSpool" from the current state, and "InSpool" is a precondition of all actions. Since "Copy Message" adds the same facts as "Move Message" (the message is filed), but does not delete "InSpool," substituting "Copy Message" for "Move Message" actions would result in a valid plan. The agent manager would prompt the user as follows:

20

"Move Message" does not apply because of the effect of a previous "Move Message" action was to move the message out of the mail spool -- you can take care of this by replacing the previous "Move Message" action by "Copy Message".
Do It

25

If no action in the plan defeated an unsatisfied precondition, the system looks for actions that satisfy the precondition and do not defeat any other actions in the plan. Consider the rule:

30

messageType = Message
=> Play Greeting "Greeting 1".

35

"Play Greeting" applies only to telephone calls, and only after a call has been picked up. Since the action "Pickup-Call" results in the call being picked up, inserting it into the plan would repair the second problem. The following message would be given to the user:

40

"Play Greeting" does not apply because
1) the current message is not a Telephone Call
- you can take care of this by selecting the
Message Type Telephone Call
Do It
2) the call has not yet been picked up -- you
can take care of this by inserting the action
"Pickup Call" before "Play Greeting"
Do It

50

The algorithm for computing repairs is thus:

55

For each unsatisfied Action A
For each unsatisfied precondition P
If \exists in plan before A, and B deletes P
Then Repairs = {B' such that $B'.add \supseteq B.add \& P \supseteq B'.delete\}$
Else Repairs = {B' such that $P \in B'.add \& B'$
can apply before A & B' does not defeat any later actions}

60

The agent manager visually indicates rules with invalid plans and problematic actions in each plan. A user may access the repair dialogue associated with an action by clicking on the action. After the user selects a repair, the system applies the repair, checks the validity of the plan again to see whether the status of any other actions has changed, and updates the repair dialogue to show the new status of the plan.

11

Returning to the operation of FIG. 8, the user selects a repair at step 512, and the agent manager applies the selected repair at step 514. The repaired plan is analyzed to ensure its validity at step 515, using rule analysis module 400 described previously. If the repaired plan is valid at step 518, the current state for the repaired plan is updated at step 520. If the repaired plan is not valid at step 518, the agent manager generates another repair.

FIG. 9 is a flow diagram of a second portion of a rule repair module used in accordance with one embodiment of the invention. If the repaired plan is valid at step 518, the agent manager checks whether the rule was a member of a problem category at step 522, and whether other rules also belong to the same category at step 524. In such cases, the agent manager asks whether the user wants to propagate the same repairs to all other rules in the category at step 526, and if so, applies the repairs to all plans within the same problem category at step 528. Once all the repaired plans are checked at step 530, the agent manager checks if all plans are checked at step 504 (shown in FIG. 8), and if so, returns. If not, the agent manager repairs the next plan. If all repaired plans are not checked at step 530, the agent manager processes the next repaired plan.

It can be appreciated by one skilled in the art that any software programs, schemes, steps or algorithms described herein with reference to an embodiment of the invention can be implemented in either hardware or software. If implemented in software, these software programs, schemes, steps or algorithms can be implemented on any general purpose computer, such as a personal computer, or special purpose computer, such as a digital signal processor (DSP), having at a minimum a processor with sufficient processing speed, and a computer readable medium such as read-only memory (ROM), programmable ROM (PROM), erasable programmable ROM (EPROM), random access memory (RAM), hard drive, optical drive, floppy disk, and so forth.

Conclusion

Although a preferred embodiment is specifically illustrated and described herein, it will be appreciated that modifications and variations of the present invention are covered by the above teachings and within the purview of the appended claims without departing from the spirit and intended scope of the invention. For example, although the rule hierarchy is defined using CLASSIC, it can be appreciated that any mechanism for creating structured descriptions of sets of objects and individual objects falls within the scope of the invention. Another example is the rule editor used by rule edit module 200, and shown in FIG. 3. It can be appreciated that any editor allowing the definition of conditions, attributes and actions, or combinations thereof, falls within the scope of the invention.

What is claimed is:

1. A method for a user to program a personal software agent using an agent manager, wherein said agent manager is connected to an input device for receiving instructions from the user, comprising the steps of:
 - creating a rule to control the personal software agent;
 - placing said rule in a hierarchical order of rules to control the personal software agent;
 - determining whether said rule conflicts with another rule within said hierarchical order; and
 - suggesting repairs to said rule if said rule conflicts with another rule.
2. The method described in claim 1, wherein the user creates said rule by defining conditions, attributes and actions of said rule.
3. The method described in claim 2, wherein said hierarchical order is comprised of parent rules, child rules, sibling rules, and inter rules.

12

4. The method described in claim 3, wherein said rule is placed in said hierarchical order using CLASSIC.

5. A method for a user to program a personal software agent using an agent manager, wherein said agent manager is connected to an input device for receiving instructions from the user, comprising the steps of:

creating a rule by defining said rule's conditions, attributes and actions;

placing said rule in an hierarchical order comprised of parent rules, child rules, sibling rules, and inter rules, wherein said rule inherits actions from said parent rules, wherein said child rules inherit actions from said rule, and said inter rules inherit actions from the intersection of said rule and said sibling rules;

determining whether said rule is valid within said hierarchical order; and

suggesting repairs to said rule if said rule is invalid.

6. The method described in claim 5, wherein said rule is placed in said hierarchical order using CLASSIC and CLASSIC determines subsumption relations among rules.

7. The method described in claim 6, wherein CLASSIC determines intersections among rules.

8. The method described in claim 7, wherein said action is defined in terms of preconditions, add list and delete list for said action.

9. The method described in claim 8, wherein said rule is repaired using an interactive dialogue explaining why said action does not apply and suggesting repairs to said rule that will enable said actions.

10. The method described in claim 9, wherein said invalid rule is categorized according to why said rule is invalid.

11. The method described in claim 10, wherein repairs applied to said invalid rule are applied to rules in the same category as said invalid rule.

12. An agent manager for a personal software agent, the agent manager having an input device for receiving instructions from a user, comprising:

a rule edit module having a rule edit input coupled to the input device for receiving instructions from the user for creating a rule to control the personal software agent and having a rule edit output;

a rule index module having a rule index input coupled to said rule edit output for receiving said rule and placing said rule in a hierarchical order of rules to control the personal software agent, and having a rule index output;

a rule analysis module having a rule analysis input coupled to said rule index output for receiving said hierarchical order and using said hierarchical order for determining whether said rule conflicts with another rule in the hierarchical order, and having a rule analysis output; and

a rule repair module having a rule repair input coupled to said rule analysis output to receive said rule if said rule conflicts with another rule in the hierarchical order in order to guide the user in selecting and applying repairs to avoid the conflict.

13. The agent manager described in claim 12, wherein said rule edit module allows the user to enter conditions, attributes and actions to define said rule.

14. The agent manager described in claim 13, wherein said hierarchical order is comprised of parent rules, child rules, sibling rules, and inter rules.

15. The agent manager described in claim 14, wherein said rule index module permits the user to create structured descriptions of sets of objects and individual objects.

16. The agent manager described in claim 15, wherein said rule index module represents said rule using CLASSIC.

17. An agent manager for a personal software agent, the agent manager having an input device for receiving instructions from a user, comprising:

13

a rule edit module having a rule edit input coupled to the input device for receiving instructions from the user for creating a rule by entering conditions, attributes and actions to define said rule, and having a rule edit output; a rule index module having a rule index input coupled to said rule edit output for receiving said rule and placing said rule in a hierarchical order comprised of parent rules, child rules, sibling rules, and inter rules to create structured descriptions of sets of objects and individual object, and having a rule index output and wherein said rule index module inherits actions from said parent rules to said rule, inherits actions from said rule to said child rules, and inherits actions from the intersection of said rule and said sibling rules to said inter rules; a rule analysis module having a rule analysis input coupled to said rule index output for receiving said hierarchical order and using said hierarchical order for determining whether said rule is valid, and having a rule analysis output; and a rule repair module having a rule repair input coupled to said rule analysis output to receive said rule if said rule is invalid in order to guide the user in selecting and applying repairs to make the rule valid.

18. The agent manager described in claim 17, wherein said rule index module determines subsumption relations among rules.

19. The agent manager described in claim 18, wherein said rule index module determines intersections among rules.

20. The agent manager described in claim 19, wherein said action is defined in terms of preconditions, add list and delete list for said action.

21. The agent manager described in claim 20, wherein said rule repair module creates an interactive dialogue explaining why said action does not apply and suggests repairs to said rule that will enable said actions.

22. The agent manager described in claim 21, wherein said rule repair module categorizes invalid plans according to why said plans are invalid.

23. The agent manager described in claim 22, wherein said rule repair module applies said repairs to rules in the same category as said invalid rule.

24. Apparatus for automatically verifying whether a new rule which is to be added to a set of rules to control a personal software agent is valid with respect to the set of rules, each rule specifying a set of conditions and a sequence of actions being interpreted in a system which causes the actions specified in the rule to be performed when the conditions specified in the rule are satisfied, the apparatus comprising:

a stored subsumption hierarchy of the rules in the set of rules; means for placing the new rule to control the personal software agent in the subsumption hierarchy; and means for using the subsumption hierarchy which includes the new rule to determine whether the new rule conflicts with another rule in the hierarchy and provide an indication when a conflict exists, wherein the means for using the subsumption hierarchy further uses the subsumption hierarchy to determine a suggested correction for the new rule when the new rule conflicts with another rule and provide the suggested correction.

25. The apparatus described in claim 24, further comprising means responsive to an input from a user of the apparatus indicating acceptance of the suggested correction for correcting the new rule according to the suggested correction.

26. Apparatus for automatically verifying whether a new rule which is to be added to a set of rules is valid with respect

14

to the set of rules, each rule specifying a set of conditions and a sequence of actions being interpreted in a system which causes the actions specified in the rule to be performed when the conditions specified in the rule are satisfied, the apparatus comprising:

a stored subsumption hierarchy of the rules in the set of rules;

means for placing the new rule in the subsumption hierarchy;

means for using the subsumption hierarchy which includes the new rule to determine whether the new rule is valid, and provide an indication of invalidity when the new rule is not valid and further using the subsumption hierarchy to determine a suggested correction for the new rule when the new rule is not valid and provide the suggested correction, wherein the means for using the subsumption hierarchy determines whether the new rule is valid by using the subsumption hierarchy to determine whether the conditions of the new rule and another rule of the set of rules can apply simultaneously and if the conditions do so apply, analyzing the actions of the rules for conflicts; and

comprising means responsive to an input from a user of the apparatus indicating acceptance of the suggested correction for correcting the new rule according to the suggested correction.

27. The apparatus described in claim 26, wherein the means for using the subsumption hierarchy determines whether the conditions of the new rule and another rule can apply simultaneously by using the subsumption hierarchy to determine whether the condition of one of the rules is more general than the condition of the other or whether the condition of one of the rules intersects with the condition of the other of the rules.

28. The apparatus described in claim 27, wherein the means for using the subsumption hierarchy further uses the subsumption hierarchy to determine a category for a suggested correction which has been accepted and find other rules which require corrections belonging to the category.

29. The apparatus described in claim 28, wherein the means for placing the new rule in the subsumption hierarchy does so each time the new rule is altered.

30. Interactive rule editing apparatus for editing a rule which is to be added to a set thereof, each rule specifying a condition and an action and being interpreted in a system which causes the action specified in the rule to be performed when the condition specified in the rule is satisfied, the interactive rule editing apparatus comprising:

input means;

output means;

means for determining whether the rule to be added to control the personal software agent conflicts with another rule with regard to the set thereof and if the rule to be added does conflict, providing a suggested correction;

means for altering the rule; and wherein the input means receives the rule to be added and provides the rule to be added to the means for determining;

the output means receives the suggested correction and outputs the suggested correction;

the input means receives an indication that the suggested correction has been accepted; and the means for altering the rule responds to the indication by altering the rule to be added in accordance with the suggested correction.



US005922045A

United States Patent [19]

Hanson

[11] Patent Number: 5,922,045
[45] Date of Patent: Jul. 13, 1999

[54] METHOD AND APPARATUS FOR PROVIDING BOOKMARKS WHEN LISTENING TO PREVIOUSLY RECORDED AUDIO PROGRAMS

[75] Inventor: Bruce Lowell Hanson, Little Silver, N.J.

[73] Assignee: AT&T Corp., Middletown, N.J.

[21] Appl. No.: 08/682,034

[22] Filed: Jul. 16, 1996

[51] Int. Cl. 6 H04M 3/56; H04M 3/50

[52] U.S. Cl. 709/206; 709/203; 709/217; 709/227; 709/228

[58] Field of Search 395/200.33, 200.36, 395/200.42-200.48, 200.57-200.59; 348/7, 12, 13; 379/93, 88, 202; 455/4.2, 5.1; 709/202-203, 205, 212-219, 227-229, 237

[56] References Cited

U.S. PATENT DOCUMENTS

5,668,863	9/1997	Bieselin et al.	379/202
5,732,216	3/1998	Logan et al.	709/46
5,740,304	4/1998	Katsuyama et al.	386/46
5,841,979	11/1998	Schulhof et al.	709/237

Primary Examiner—Zarni Maung

Assistant Examiner—Bharat Barot

[57] ABSTRACT

An audio service monitors the progress of a user as he or she proceeds through the service. The location in the service at the time of termination is detected and correlated with user and service identification information. As a result, a bookmark is created which enables the user to later resume the service at the point of prior termination.

30 Claims, 3 Drawing Sheets




FIG. 1





FIG. 2A



COPY

ORIGINAL
FILED

1 BAKER BOTTS L.L.P.
2 Bryant C. Boren Jr. (SBN 260602)
3 Kevin E. Cadwell (SBN 255794)
4 620 Hansen Way
5 Palo Alto, CA 94304
6 Telephone: (650) 739-7500
7 Facsimile : (650) 739-7699
8 Email: bryant.c.boren@bakerbotts.com
9 Email: kevin.cadwell@bakerbotts.com

2010 MAR 12 A 10:23

ADR

RICHARD W. WICKING
CLERK, U.S. DISTRICT COURT
NORTHERN DISTRICT OF CALIFORNIA

6 Attorneys for Plaintiffs
7 AT&T Intellectual Property I, L.P. and
8 AT&T Intellectual Property II, L.P.

E-filing

5809,492 -
5,922,045 -
6,118,976 -
6,983,478

9
10 IN THE UNITED STATES DISTRICT COURT
11 FOR THE NORTHERN DISTRICT OF CALIFORNIA
12 SAN JOSE DIVISION

13 AT&T INTELLECTUAL PROPERTY I, L.P., and
14 AT&T INTELLECTUAL PROPERTY II, L.P.,

CASE NO. 10-cv-01343

MEJ

15 Plaintiffs,

**COMPLAINT FOR PATENT
INFRINGEMENT**

v.

16 TiVo INC.,

17 Defendant.

DEMAND FOR JURY TRIAL


18
19 Plaintiffs AT&T Intellectual Property I, L.P., and AT&T Intellectual Property II, L.P.
20 (collectively, "AT&T"), by its undersigned attorneys, allege in this Complaint against TiVo Inc.
21 ("TiVo") as follows:

NATURE OF THE ACTION

22 1. This civil action arises out of patent infringement under the Patent Laws of the
23 United States, 35 U.S.C. § 100, *et seq.*, and, in particular, under the provisions of
24 35 U.S.C. § 271.

25 2. AT&T seeks relief from TiVo's infringement of AT&T's patent rights under U.S.
26 Patent Nos. 5,809,492; 5,922,045; 6,118,976; and 6,983,478, as set forth more fully below.

FIG. 2B



**METHOD AND APPARATUS FOR
PROVIDING BOOKMARKS WHEN
LISTENING TO PREVIOUSLY RECORDED
AUDIO PROGRAMS**

BACKGROUND OF THE INVENTION

The present invention is directed to providing a bookmark for audio programs. In one embodiment an audio service includes the capability of creating a user bookmark to enable a user to proceed through an audio program at that user's own pace.

It is known to provide telephone-based audio services over, for example, the Public Switched Telephone Network (PSTN). For example, it is known to provide call up services where a service user dials a particular number through the PSTN to obtain information. One example of such a service is a "900" service that enables a user to call a "900" number and receive the day-to-day comments of a celebrity or sports personality. Similarly, it is possible to provide an audio service in which the user will dial a designated number and have audio program information, such as news summaries, played back to the user through the PSTN.

The presently available audio services are somewhat limited in nature. In particular, audio-based services that provide serial information are not as powerful as they could be because the user is not provided with a way to select the appropriate pace with which they will proceed through the material. For example, there presently is no flexibility provided to the user for selecting a start point, or more importantly a restarting point if the user accesses a given audio service multiple times. In a typical service, if the user chooses to disconnect from the service at any point prior to the completion of the program material, the call terminates and the service takes no note of where or when the user terminated the access to the service. Connect time may be monitored for billing purposes, but no correlation is drawn to the user's progress through the audio program. If the user later reconnects to the service, the program material is cued to its initial starting point and the user must listen to the entire portion of the program material which was accessed earlier before returning to the point at which the service had earlier been terminated.

The present construction of these services limits their usefulness in providing audio services, such as audio books or audio "soap operas", since there is no possibility of quickly returning to the point at which the user had previously terminated the service. The user is discouraged from re-accessing such services.

For example, the user could with some difficulty return to the same part of a story or soap opera, at the cost of having to listen to the same program material again before proceeding on to new material. Thus, it would be advantageous if there were some technique provided for allowing a user to enjoy the audio services at a pace set by the user.

A similar problem arises in the context of audio program material that can be downloaded from a network audio service, e.g., a music program could be accessed via the Internet and downloaded to a user's PC. Upon playback of the program the user may decide to stop playing the program to pursue some other interest. It would also be advantageous if the PC could provide a technique for automatically positioning the playing back of the downloaded program to avoid having to search for the desired re-starting point.

SUMMARY OF THE INVENTION

The present invention provides a "bookmark" that permits a user who is reviewing audio program material to stop at

any point in the program and to resume the review at the point at which the user previously stopped.

In accordance with an embodiment of the present invention, a user can access an audio information services platform through the PSTN. The platform has an audio content (program) database and a personal profile database. The user has an assigned user code. When the user connects to the audio information service platform, the user code is utilized to access information stored in the personal profile database. The stored information identifies the services previously accessed by the user and identifies the location within each audio service at which the user has previously terminated the service. The user can then select a desired audio service and request to be returned to the location in the audio service at which the previous access had been terminated.

In accordance with the embodiment the bookmark is generated by monitoring or tracking the progress of the user in reviewing the program material. In particular, a playback module may be coupled to a position location detector such as a block counter or timer. When the user terminates the access, the user ID, a service identifier and a last position location are grouped and stored in the personal profile database for later use.

By providing a bookmark for the audio services the user can better access and utilize audio books, multi-step programs such as stop-smoking programs, audio soap operas, travel direction services, or any serial presentation of information that lends itself to self-pacing by the user.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 discloses an audio information service system in accordance with an embodiment of the present invention.

FIG. 2A illustrates a flowchart for creating a bookmark in the audio service system of FIG. 1 in accordance with an embodiment of the present invention.

FIG. 2B illustrates a flowchart for using a bookmark created in FIG. 2A.

DETAILED DESCRIPTION

FIG. 1 illustrates, in block diagram form, a system for providing an audio information service that includes bookmarks. The service provider has an audio information service platform 103 which is accessible through a Public Switch Telephone Network (PSTN) 102. Typically in such circumstances the platform is accessible by dialing a particular number, e.g., 800-XXX-YYYY, that is associated with the platform. The user can access the platform through the PSTN simply using the telephone 101. The telephone will include a keypad which enables the user to enter alpha-numeric codes represented by DTMF signals that can be transmitted through the PSTN 102 to the audio information service platform 103. One known audio services platform that does not include the bookmark capability is the CONVERSANT® system, belonging to Lucent Technologies, Inc.

In the present invention the platform 103 includes not only a database for audio content, 104 but a personal profile database 105 that stores information about the users of the service. A service controller 107 is a processing device which has the capability of controlling accesses to the personal profile database to obtain information about the users that are accessing the service platform. Alternatively, the service controller can arrange to have information about a user stored in their personal profile database 105.

The service controller 107 also interfaces with a program retrieval and playback module 109. This module is responsible for accessing audio from the audio content database 104 under the specific instructions of the service controller. For instance, the service controller will provide information to the module as to which audio content to obtain as well as identify the location within the audio content at which the service should begin. Typically each service will have an associated identifier or code that uniquely identifies the service and/or content. The service controller then arranges for the playback audio service to be routed to the appropriate user through the PSTN 102 to the user's telephone 101.

The service controller also interfaces with a service menu module 108 that provides information about the available audio services in the audio content database. Under the control of the combination of the service controller and the service menu module the user, upon connection to the platform, is given information about the available services and, as is described below, can be instructed about options for resuming previously accessed audio services.

The service controller could also interface with an Automatic Speech Recognition Unit (ASR) which operates to detect voice responses by the users to menu prompts, rather than detecting keypad or DTMF responses.

The service controller 107 also interfaces with a user ID module 106. This connection provides the service controller with information about the user including the identity of the user for use together with the personal profile database. User IDs are well known. One possibility for a user ID is the detection of the automatic number information (ANI) which can be transmitted by the PSTN 102 and which identifies the telephone number associated with the line connected to the user's telephone. Alternatively, the user can be prompted to provide a personal ID or a PIN so that the user can access the audio service from any location rather than be limited to accessing it from one particular telephone.

Finally, the audio information service platform 103 also includes a position detector, shown as a block counter in this embodiment, which is used in conjunction with the playback module 109 to monitor the user's progression through the audio service. The block counter provides useful information about where the user is in a selected audio service at any given time. This present position location information is then detected and associated with a user when the user terminates an audio service. The service controller can then provide the user information (user ID), service identification information (service ID) and termination location information to the personal profile database at the time the user terminates a service.

The entire audio service platform is operated with the service controller 107 operating under the control of a stored program for accessing the audio services, creating a bookmark and using an existing bookmark. The accessing of audio services itself is well known. However, the creation of a bookmark and the use of an existing bookmark is described below with reference to the flowcharts in FIGS. 2A and 2B.

One potential embodiment for creating a bookmark is described by the flowchart of FIG. 2A. In accordance with this operation, a customer first accesses the audio information service platform through the PSTN and the platform identifies the user either by detecting the ANI or by prompting the user to enter a PIN, step 201. The user then is provided with an interactive service menu and selects an audio program from the menu, step 202. The interactive service menu can operate in conjunction with an ASR and/or a DTMF detector to detect the user's responses to the menu

prompts. The service menu module reports the selection of the audio service to the service controller 107 and in turn, the controller activates the playback module to request a playback of the desired audio service, step 203. The program retrieval and playback module 109 then acts upon the request from the service controller to access the audio content database 104 to select the desired audio service in accordance with the service ID associated with the selected service. The playback module 109 then can check the selected program for length, step 204. This can be done for instance by checking the file length or length information could be encoded into a header, for example. It may be determined in advance that certain programs are too short to warrant providing bookmarks to the user so that only selected programs will provide the bookmark capability. If that is the case, then having obtained the information identifying the length of the program, the playback module determines whether the program is long enough to warrant bookmarks, step 205. If the module detects that the service is not long enough to warrant bookmarks then the module will playback the program in the normal way without tracking the user's progress through the program, step 206.

If, however, the playback module detects that the program is long enough to warrant bookmarks, the playback module fetches the first block of the audio program and records the block number in the block counter 110 in step 207. As playback of a given block of the audio program is completed, the playback module fetches the next block in the sequence and begins to play it while updating the content of the block counter to ensure that the counter maintains an accurate indication of the block now involved in the playback process, step 208. Simultaneously, the service controller continuously monitors for a user request to end the program, step 209. The user can either terminate the program by selection of specified keys on the user telephone keypad or can simply hang up. If the user is still listening to the program as detected in step 210, then the operation recycles through steps 208 and 209 to ensure that the user continues to get additional blocks of the audio service and that the location of the user through the service is constantly monitored. If, however, the user is no longer listening to the program, then the service controller notifies the playback module to terminate the playback, step 211. At the time of termination the playback module passes the last recorded block number stored in the block counter as well as the program ID for the service in use to the service controller, step 212. The service controller then takes the location information and the program information and combines it with the customer or user ID which is detected in step 201 and records this information in the personal profile database 105 for future reference, step 213.

The creation of the bookmark is variable depending upon the type of program that is being provided to the user. Different types of programs lend themselves to different ways to monitor the user's progress through the program. For example, in a book environment it may be desirable to keep track of the user's progress through the book in terms of blocks that relate to either chapters, pages, or paragraphs. Then, as each element, for example paragraph, is accessed, the block counter is increased by one so that the system keeps track of which paragraph in the presentation the user is presently located. Alternatively, the program may be a multi-step improvement program, e.g., a stop-smoking program. In such a circumstance, it may be beneficial to detect the user's progress through the individual steps. In that case the blocks referred to in connection with FIG. 2A may in fact, correspond to the steps (or sub-steps) through the

process or program. Additionally, it may be advantageous to detect the actual elapsed time of the audio program. This would be applicable to the actual playback of a continuous program where the system monitors the time (the number of minutes and seconds) that the user has progressed into the audio file. Therefore, the block counter 110 could be replaced by a timer and the timing information associated with the location of the user through the audio content could be then provided to the service controller at the time of termination. Then the timing information would be stored with the program ID, and the user ID in the personal profile database.

FIG. 2B illustrates a flowchart for use of the bookmark in connection with providing the audio service.

At the starting point the customer dials the audio information service via the PSTN, step 221. The system answers the user's call and the service controller 107 activates the user ID module to detect the identity of the user, step 222. As described above, the user module can either use ANI or the PIN specifically entered by a user to identify the user, step 223. The service controller then queries the personal profile database to verify the user ID and to determine whether any bookmarks exist and to determine, based on past experience with the user, whether the user has any particular audio service preferences, step 224. The controller then detects whether any bookmarks exist for the user in decision step 225. If no bookmarks exist for the user then the service controller calls the service menu module with a standard non-bookmark menu similar to the menu which was referred to above in connection with FIG. 2A, step 226. That menu provides the user with an identification of the audio information services that are available through that platform. The service module then plays the standard menu, step 227 and the process continues from step 203 of FIG. 2A.

If, however, the service controller detects that a bookmark has been created for this user then the service controller retrieves the program or programs with bookmarks and the specific bookmark values from the personal profile database, step 230. The service controller can then pass the identification information regarding the previously accessed services in which bookmarks exist to the service menu module, step 231. The service menu module then can (step 232): play a special menu that includes the programs that have bookmarks associated therewith; and, provide the user with the option of continuing the program from where the user last left off or to start the program from the beginning. In one embodiment the user indicates the choice by activating a corresponding key of the user's keypad to indicate the selection of the program and to select the location to begin playback of the program. The service menu module then indicates which selections the user has made, step 233. If the user has opted to begin a program from the beginning rather than from the place where the previous access was terminated, then the service controller instructs the playback module to play the program from the beginning, step 234, and the process continues from step 207 of FIG. 2A. If the user has indicated a desire to continue the program from where they last left off, the service controller passes the bookmark value to the playback module with the instruction to play the program from the bookmark, step 235. The playback module uses the bookmark value to index a point in the program where the user left off. The playback module then requests the next block of program from the audio content database, step 236. As indicated above, if the increments of the program are in time increments or step increments, the playback module would then proceed to

mark the restart location according to the time or step information provided with the bookmark value. The playback module then begins to play the program from the bookmark and starts the block tracking or location monitoring for bookmark updates by carrying out the process from step 208 of FIG. 2A onward, step 237.

Of course, if the user selects a new service for which a bookmark has not yet been created, the operation of the system would continue from step 203 of FIG. 2A.

The above block diagrams and flowcharts show one embodiment for a system that provides audio services and creates and uses bookmarks in connection with providing those audio services. It should be recognized that various modifications to the embodiment are possible. For instance, it is possible that the service will be a revenue generating service. In such a circumstance the user may be prompted to enter information not only identifying the user but also may be prompted to provide information for billing purposes. This billing information can also be stored in the personal profile database. Of course, the user ID and billing information codes can be one and the same. As indicated above, the user may have the capability of indicating user preferences for types of audio services. This preference information can also be stored in the personal profile database and used to construct the appropriate menu to be provided by the service menu module. Similarly, it should be noted that the present invention is not limited to user telephones over a PSTN. Any kind of network that will support audio services can be an appropriate transmission medium for the audio service from the platform to the user (for example, a wireless network). Similarly, the user may employ a PC or other device or personal appliance to access the audio information service with, e.g., appropriate user ID information. As an example of an alternative application, a user may access an audio program or service via the Internet. In such a circumstance, available bandwidth may limit the viability of real-time audio program presentation. Thus, the more appropriate presentation would involve first downloading the program material to the user's PC where it would be stored for playback. Once the user begins to playback the program the PC could monitor the progress through the program material and generate the appropriate bookmark in a manner consistent with the techniques described above. Then when the user terminates the program and subsequently returns they will be given the option of where to re-start the program. The termination and creation of a bookmark are distinct from a pause functionality in that the termination actually releases the program and ends its execution. The pause function keeps the program indefinitely cued. Thus, the bookmark would be useful in this environment as well.

All of these capabilities are ancillary to the invention and are different techniques for employing that invention. The invention provides the capability of the user to proceed through an audio program at the user's own established pace and to return to the program and to the last location in the service at will. As a consequence, it provides the capability for creating and utilizing audio services that contain serial or continuous information in a more user friendly manner.

What is claimed is:

1. A method of facilitating a user's review of previously recorded audio program material over at least two review sessions, the audio program material having been communicated to the user over a communication network, the method comprising the steps of:
monitoring a user's progress in the user's review of audio program material during a first review session; and
when the user terminates the first review session, storing in a memory associated with the audio program mate-

rial an indication of the user's progress in reviewing the audio program material during said first review session, wherein said stored indication can be utilized in a second review session subsequent to the termination of said first review session.

2. The method of claim 1, further comprising the step of an audio program service furnishing said audio program material to said user over a communication channel.

3. The method of claim 2, wherein said communication channel comprises a data network.

4. The method of claim 2, wherein said communication channel comprises a telephone network.

5. The method of claim 2, wherein said communication channel comprises a wireless communication channel.

6. The method of claim 1, wherein said steps are implemented by a user's personal appliance.

7. The method of claim 6, wherein said audio program material is stored on CD-ROM.

8. The method of claim 1, wherein said steps are implemented by a communications network-based service.

9. A method of facilitating a user's review of previously recorded audio program material over at least two review sessions, the audio program material having been communicated to the user over a communication network, the method comprising the steps of:

monitoring a user's progress in the user's review of audio program material, said monitoring performed during a first audio program review session;

when the user terminates the first audio program review session, storing in a memory associated with the audio program material an indication of the user's progress in reviewing the audio program material during said first audio program review session;

in a second audio program review session subsequent to the termination of said first audio program review session, playing said audio program material to said user beginning from a position within said material based on said stored indication.

10. The method of claim 9, further comprising the step of an audio program service furnishing said audio program material to said user over a communication channel.

11. The method of claim 10, wherein said communication channel comprises a data network.

12. The method of claim 10, wherein said communication channel comprises a telephone network.

13. The method of claim 10, wherein said communication channel comprises a wireless communication channel.

14. The method of claim 9, wherein said steps are implemented by a user's personal appliance.

15. The method of claim 14, wherein said audio program material is stored on CD-ROM.

16. The method of claim 9, wherein said steps are implemented by a communications network-based service.

17. A method of creating a bookmark for use with an audio service that provides previously recorded sequential audio information comprising the steps of:

associating a user code with a user of the audio service; monitoring a present location of the user in a sequence of audio information in an access to the audio service; detecting a termination of the access to the audio service; creating a termination code defining the present location of the user in said sequence at the time the termination is detected;

correlating the user code and the terminating code; and storing the result of said correlating step in a memory associated with the audio service for later use.

18. The method of claim 17, wherein a service identifying code is associated with each audio service; and said step of correlating further includes the step of associating a service identifying code with said user code and said termination code.

19. The method of claim 17, wherein said step of monitoring comprises the steps of loading a register with an initialization value at the start of the service and updating the register as the audio information is presented.

20. The method of claim 19, wherein said audio information is divided into discrete blocks of information and wherein said step of updating includes the step of revising the contents of said register to hold a block identifier corresponding to the block of information being conveyed by the service at that time.

21. The method of claim 19, wherein said register stores the time elapsed from the beginning of the providing of the audio information.

22. A method of providing audio services using a bookmark comprising the steps of:

generating a menu of a plurality of audio services; detecting a selection of one of said plurality of audio services; transmitting previously recorded audio information from the selected audio service; monitoring a user's position in the selected audio service as the corresponding audio information is transmitted; detecting a termination of the selected audio service that occurs prior to completion of said service; creating and storing in a memory associated with at least one of said plurality of audio services a bookmark that identifies a user, the selected service and the user's position in the selected service at the time of termination; subsequently accessing said bookmark; and returning the user to the location of the selected service based on said bookmark.

23. The method of claim 22, wherein said selected audio service comprises a plurality of discrete blocks of audio information, each block having a unique block identifier wherein said step of tracking comprises the step of temporarily storing the block identifier of a discrete block of audio information as that information is transmitted.

24. The method of claim 22, wherein said step of monitoring comprises the step of monitoring an elapsed time from a time at which said transmitting step begins.

25. An audio information service platform for providing previously recorded audio information over at least two review sessions comprising:

an audio content database; a personal profile database; a program playback module coupled to said audio content database; a playback position monitor coupled to said program playback module; and a service controller utilizing a user identifier, and storing an audio content identifier and a playback position identifier with said user identifier in said personal profile database in response to termination of a first review session, wherein said playback position identifier indicates a user's progress in reviewing audio information during said first review session and can be utilized in a second review session subsequent to the termination of said first review session.

26. The platform of claim 25, wherein said audio content database stores an audio information for a plurality of audio services.

27. The platform of claim 26, further comprising a service menu module, coupled to said service controller and identifying an audio service in said audio content database that is desired by a user.

28. The platform of claim 27, wherein said service menu module advises said service controller of a desired starting point for an identified audio service.

29. The platform of claim 25, further comprising a user ID module that identifies a service user and wherein said

service controller searches said personal profile database for data relating to an identified service user.

30. The platform of claim 29, wherein said service controller is coupled to said program playback module to initiate a resumption of an audio program at a location defined by a playback position identifier associated with an identified service user in said personal profile database.

* * * * *

EXHIBIT C



US006118976A

United States Patent [19]

Arias et al.

[11] Patent Number: **6,118,976**
 [45] Date of Patent: **Sep. 12, 2000**

[54] **ASYMMETRIC DATA COMMUNICATIONS SYSTEM**

[75] Inventors: **Salvador Luis Arias, Chamblee; Edward Irby Comer, Marietta; Roy Curtis Dunn, Douglasville; Melvin Duane Frerking, Norcross; Fred Thomas Danner, III, Alpharetta; Richard Sammis Bergen, Jr., Roswell; Sidney Walker Elliott, Atlanta; Thomas Franklin Evans, Stone Mountain; Craig Brent Chambers, Stockbridge, all of Ga.**

[73] Assignee: **BellSouth Intellectual Property Corporation, Wilmington, Del.**

[21] Appl. No.: **08/994,531**

[22] Filed: **Dec. 19, 1997**

Related U.S. Application Data

[63] Continuation of application No. 08/447,537, May 24, 1995, abandoned.

[51] Int. Cl.⁷ **H04N 7/173**

[52] U.S. Cl. **455/5.1; 348/12; 348/7**

[58] Field of Search **345/327; 348/12, 348/13, 7, 6; 455/4.2, 5.1, 3.1, 6.1; H04N 7/16, 7/173**

[56] **References Cited**

U.S. PATENT DOCUMENTS

4,747,160 5/1988 Bossard.
 5,563,892 10/1996 Kostreski et al. 370/69


5,600,364 2/1997 Hendricks et al. 348/1
 5,668,610 9/1997 Bossard et al.

*Primary Examiner—Chris Grant
 Attorney, Agent, or Firm—James L. Ewing, IV; Geoff L. Sutcliffe; Kilpatrick Stockton LLP*

[57] **ABSTRACT**

An asymmetrical data communications system (ADCS) provides point-to-multipoint television programming including conventional television programming, near video-on-demand (NVOD) or video-on-demand (VOD), and the full variety of available programming, via a compressed, digitized UHF transmission. A program subsystem of the ADCS receives programming from content providers and processes the received signals for channel and VOD or NVOD service, then sends the aggregated signal to a transmission subsystem that modulates, channelizes, amplifies, filters and broadcasts the digital UHF signals over the air. Subscribing viewers are equipped with an intelligent control box (ICB) suitably configured to receive, demodulate, and decode the digital UHF broadcast and to transmit the resulting signal to one or more display or other terminal devices. The ICB further provides a matrix switch or gateway for receiving signals over any available transmission path. In addition to the ICBs of the subscribing viewers, a return path subsystem of the ADCS system includes a session control and administrative facility to which the ICBs are linked via the public switched telephone network or suitable wireless alternative so that transaction and viewing data can be received from the subscribing viewers. The return path subsystem, in turn, is linked to the program subsystem in order to route to that subsystem any information necessary or useful for providing programming.

31 Claims, 9 Drawing Sheets







FIG 1



THE PARTIES

3. Plaintiffs AT&T Intellectual Property I, L.P. and AT&T Intellectual Property II, L.P. are Nevada limited partnerships with headquarters in Atlanta, Georgia.

4. Defendant TiVo Inc. is a Delaware corporation with a principal place of business at 2160 Gold Street, Alviso, California 95002.


JURISDICTION

5. This Court has subject matter jurisdiction over this dispute pursuant to 28 U.S.C. §§ 1331 and 1338(a).

6. This Court has personal jurisdiction over Defendant TiVo, which has conducted and continues to conduct business in the State of California and in this Judicial District. Defendant TiVo's principal place of business is within this Judicial District, in Alviso, California. Defendant TiVo has committed acts of patent infringement alleged herein within the State of California and, more particularly, within this Judicial District. Moreover, Defendant TiVo has purposefully and voluntarily placed its infringing products, processes, and services into the stream of commerce with the expectation that they will be purchased by consumers in this Judicial District. These infringing products, processes, and services have been and continue to be purchased by consumers in this Judicial District.

VENUE

7. Venue is proper in the Northern District of California under 28 U.S.C. §§ 1391(b) and 1400(b) because, upon information and belief, acts and transactions constituting at least a subset of the violations alleged herein occurred in part in this Judicial District and because Defendant TiVo is found and transacts business in this Judicial District. Venue is also proper in this Judicial District under 28 U.S.C. § 1391(c) because Defendant TiVo is subject to personal jurisdiction in this District.

**FIG 3**




FIG 4






FIG 5








FIG 8




FIG 9

ASYMMETRIC DATA COMMUNICATIONS SYSTEM

RELATED APPLICATIONS

This application is a continuation of Ser. No. 08/447,537, filed May 24, 1995, now abandoned.

FIELD OF THE INVENTION

The present invention relates in general to the field of data communications, and in particular, to the fields of television and telecommunications.

BACKGROUND OF THE INVENTION

Originally, television programming was provided solely by over-the-air broadcast. The most widely received over-the-air transmissions were and continue to be in the very high frequency (VHF) band and only to a much more limited extent in the ultra high frequency (UHF) band of the RF spectrum. In recent decades, however, the delivery of television programming into the home increasingly has been delivered via CATV (cable) service. CATV transmission offered considerably higher bandwidth than was available over-the-air, while the quality of its transmission—for those equipped with the necessary coaxial cable for receiving the analog signal, and the hardware required for descrambling it—has been generally better than analog airwave transmission, which is subject to a variety of forms of signal interference.

The high bandwidth and transmitted signal quality of CATV transmission relative to over-the-air broadcasting has led to CATV being a dominant force in the market for multiple channel programming. That CATV coaxial delivery systems in principle provide sufficient bandwidth to permit two way communication with subscribers has fueled speculation that CATV may provide an early venue for the provision of interactive television services. CATV service, however, is inherently limited by the extent of its geographic penetration. CATV service is simply unavailable in locations that cable providers have chosen not to serve. Even where CATV service is available, installation of the coaxial cable is disruptive as well as expensive.

The provision of interactive programming content has become a major goal of the converging television, entertainment, programming, telecommunications, on line and computing industries. Telecommunications companies, for example, have invested significant resources in developing "video dialtone," in which programming services would be provided by way of the telephone system. Although this approach shows enormous potential and the ability to revolutionize both the telecommunications and television industries, it presents certain non-trivial technological and economic problems. The provision of video dialtone services to the home may turn on the installation of optical fiber and/or coaxial cable in place of existing twisted pair telephone connections, which presents a highly expensive and time consuming proposition that is not expected to be implemented for some time.

Aside from efforts to move to what has become known as high-definition television (HDTV), and to provide such HDTV services over the air, comparatively little attention has been paid recently to enhancing television programming services that are delivered over the airwaves. A possible reason for the comparative lack of effort may be that the perceived need to support interactivity would seem to militate against a video delivery system that uses as its transport

medium one that apparently lacks a return path. The available choices for delivery of television programming, meanwhile, have continued to grow into such areas as "wireless cable" and direct broadcast satellite, tending to draw attention even further away from over-the-air broadcasting.

Another factor barring or at least complicating any attempt to provide enhanced television service over-the-air is imposed by existing federal communications regulations, set forth in Title 47 of the United States Code of Federal Regulations. These regulations establish a framework for the delivery of programming over-the-air that is as rigid as it is highly detailed, and that, short of legislation, is subject to modification only by the Federal Communications Commission, and then only within the mechanism provided by the Administrative Procedures Act, 5 U.S.C. §§551-559, §§701-706, §1305, §3105, §3344, §5372, §7521.

Still another obstacle to the provision of an alternative means for providing programming services on a large scale, such as by over-the-air broadcast, has to do with equipping intended recipients of the programming services with the means to receive and view the programming. Aside from the direct expense associated with providing a new device, viewers may harbor concerns about the compatibility of a new in-home device with any existing premises equipment, in which a viewer may have made a sizable investment. Even if a device were to be provided gratis by a service provider, for example, subscribers may be somewhat reluctant to commit to a particular system if it were not compatible with existing delivery systems purely for reasons having to do with such things as clutter and the consumption of available space. Nevertheless, the apparent demand not only for traditional television programming, but also for such services as home shopping, video games, data services such as electronic catalogs, stock market quotations, sports scores, and electronic newspapers, as well as interactive services and video on demand (VoD) or near video on demand (NVoD) continues to grow. This consumer demand, coupled with an increased demand by marketing organizations for demographic and consumer preference information for use in their characterizing and targeting the increasingly segmented consumer populace faced with a growing number of viewing alternatives makes clear that any alternative means for delivering programming must have a return path for enabling viewer interaction.

SUMMARY OF THE INVENTION

The system, method and device according to the present invention solves the problems described above by providing an asymmetric data communications system (ADCS) capable of furnishing an alternative to conventional over-the-air and CATV television transmission, and that is also capable of providing functionality not furnished by either of those delivery systems. The system according to the present invention provides an alternative means for the delivery of video and audio entertainment programming, as well as a variety of data services, such as electronic catalogs, stock market quotations, sports scores, electronic newspapers, and will be able to carry services that have not yet been conceptualized. The ADCS system of the present invention at the same time provides a return path capable of supporting viewer interactivity, enabling the viewer to request transactions and orders for services that require authorization, to engage in interactive participation in programming distributed on a programming channel and other forms of interaction, as well as providing a data path and a mechanism for gathering demographic information from subscribers.

A first aspect of the system according to the present invention makes available previously unusable over-the-air broadcasting spectrum. Briefly, and as described at length below, the system according to the present invention digitizes, compresses and modulates signals for transmission in the UHF spectrum band that is currently assigned to television broadcasters, the signal compression using presently available techniques to achieve eight or more times the capacity using conventional techniques as can now be carried in that band. The ADCS system according to the present invention thereby provides point-to-multipoint multichannel broadcast services over-the-air on a scale previously available only using cable, and does so with high quality transmission and reception, and without interfering with existing channels broadcasting under the NTSC (National Television System Committee), PAL (Phase Alternation Line), SECAM (Sequential Couleur A Memoire), or other color television transmission standard. In addition, a program subsystem of the present invention aggregates heterogeneous programming from a number of content providers for digital UHF transmission, as well as VOD and NVoD, both services being provided only upon completion of an authorization function.

The system according to the present invention provides an alternative to CATV for those who wish to subscribe to services that have previously been considered the exclusive province of CATV. Notably, subscribers dissatisfied with CATV services or who simply are not served by CATV will be able to gain access to a large number of channels and at a considerably lesser expense, since the below-described system will enable program providers to offer the point-to-multipoint broadcasting at a cost lower than that associated with that of the installed cable system base. Subscribers will be able to receive and participate in interactive television without the need for a cable connection. Such subscribers will simply need a suitable device, referred to herein as an intelligent control box (ICB) and described in detail below.

According to the present invention, the ICB is adapted not only to receive, decompress, decode and transmit for display the received digitized UHF signals, but also to provide a terminal capable of establishing a return path to the broadcaster via the public switched telephone network (PSTN) (e.g., POTS, ISDN, ADSL, B-ISDN) or a suitable wireless alternative. The ICB serves as an electronic gatekeeper, providing matrix switch functionality to interface between the ADCS and a consumer television, computer or any suitable monitor or terminal device. As further described below, the ICB can include transmission decoding functionality, data storage, switching and authorization functions. The ICB can also include a capability for switching between a variety of non-ADCS inbound or downstream sources, including conventional "over-the-air" television, CATV, MMDS (multipoint microwave distribution system, or "wireless cable"), DBS ("direct broadcast satellite"), LMDS ("local multipoint distribution system," provided, e.g., by Cellular Vision), VCRs, computer/video games, and mass storage devices.

A final part of an ADCS system according to the present invention is a return path facility which, along with ICBs to which it is linked by the PSTN or wireless alternative, constitutes the return path subsystem of the ADCS system according to the present invention.

Accordingly, it is an object of the present invention to provide an alternative to conventional wired cable television (CATV).

It is another object of the present invention to provide this alternative to conventional wired cable television in the form

of an over-the-air radio frequency (RF) point-to-multipoint broadcast and receiving system.

It is a further object of the present invention to provide an alternative to conventional wired cable television that includes a return path by way of which user transactions, orders, demographics and other information may be sent or collected from a subscriber premises device.

It is yet another object of the present invention to provide the return path in the form of a wired point-to-point configuration, such as a configuration using switched (analog or digital) telephone technology, that in the context of a digital UHF point-to-multipoint broadcast system provides an asymmetrical data communications system.

It is another object of the present invention to provide an ADCS having a program subsystem capable of aggregating a variety of heterogeneous programming, digitizing and compressing these signals for UHF broadcast.

Another object of the present invention is to provide an ADCS having a return path facility accessed by subscribers over a public telecommunications network for capturing and fulfilling program requests and other transactions and also for collecting subscriber demographics information.

It is still a further object of the present invention to provide an intelligent control box to act as an electronic gate keeper at the subscriber's presence in the form of a matrix switch capable of providing an interface, capable of receiving and decoding a digital RF transmission in the UHF band, and communicating subscriber messages to the ADCS.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a schematic view of an embodiment of an asymmetrical data communications system according to the present invention.

FIG. 2 is a component block diagram of the components of an embodiment of an asymmetrical data communications system according to the present invention, showing data paths between the components.

FIG. 3 is a block diagram showing the architecture and functions of the major components of an embodiment of an asymmetrical data communications system according to the present invention.

FIG. 4 is a component block diagram of a program subsystem of an embodiment of an asymmetrical data communications system according to the present invention.

FIG. 5 is a block diagram showing the signal processing steps of multiple video sources performed in the program subsystem and/or transmission subsystem of an embodiment of an asymmetrical data communications system according to the present invention, and indicating alternative suitable techniques for performing each step.

FIG. 6 is a block and partially schematic diagram showing the transmission subsystem of an embodiment of an asymmetrical data communications system according to the present invention.

FIG. 7 is a component block diagram of an ICB of an embodiment of an asymmetrical data communications system according to the present invention.

FIG. 8 is a component block diagram of a return path facility of an embodiment of an asymmetrical data communications system according to the present invention.

FIG. 9 is a logic flow diagram for a representative operation of the ICB and return path facility of an asymmetrical data communications system according to the present invention.

DETAILED DESCRIPTION OF THE INVENTION

The system and components of the system according to the present invention will be described with reference to the accompanying figures, and in an order that reflects the delivery of services to consumers. A high-level view of an embodiment of the system according to the present invention is provided in FIG. 1. As shown, a program subsystem 10 is equipped to receive content from various sources, including: non-local content providers 12 via satellite or any other suitable means or transmission path (including optical fiber, for example); from local TV programming entities 14 via microwave, optical fiber, cable, or other suitable transmission paths; or over any suitable path from any other source (not shown). The received content may include conventional channel television broadcasts, home shopping, data services, electronic catalogs stock market quotations, sports scores, electronic newspapers and other content, whether or not presently available. Video-on-demand (VoD) or near video-on-demand (NVoD) may also be received by any known means and provided by conventional video servers, as described at greater length below.

Program subsystem 10 collects and processes the signals from these various sources and, after processing the signals, provides the programming, data and any other received content in an appropriate form and over a suitable data link 18 to a transmitter subsystem or site 20.

Transmitter subsystem 20 includes transmission equipment, described in detail below, for generating a signal having sufficient effective radiated power (ERP) and signal-to-noise ratio (SNR) to reach a set of intended subscribers 30 with acceptably high quality reception. In the preferred embodiment of the present invention, program subsystem 10 and transmitter subsystem 20 transmit this combined content in digital form and in the ultra high frequency (UHF) band of the RF spectrum (407-806 MHz).

Subscribers 30 each are equipped at their premises with appropriate receiving and processing equipment (not shown in this view, but described in detail below). Using this equipment, subscribers 30 can select from among the variety of content carried by the signal(s) transmitted by the equipment at transmission site 20, as well as from among the full array of other programming sources as to which the premises equipment acts as an electronic gateway. Communications from subscribers 30 can include requests for VoD, NVoD, or other interactive or non-interactive program service, and can also include marketing information regarding subscriber 30. Messages from (or about) subscriber 30 are transmitted via switched telephone technology, i.e., the public switched telephone network (PSTN) lines 40 (or a suitable wireless alternative, not shown) and via switching and other service sites maintained and operated by a local telephone service provider 45, to a return path facility 50 (which may also be operated by a telephone service provider, BellSouth Corp., for example).

Return path facility 50 may also undertake transactions with other service providers (collectively identified by reference numeral 48). Return path facility 50 also receives communications from other subscribers (via lines 49A, 49B, . . . , 49N), and performs a variety of functions including the control of sessions with subscribers 30 and conducting administrative functions, both of which will be discussed in detail below. Finally, return path facility 50 communicates authorization requests to program subsystem 10 via a high capacity data link 52, such as a broadband circuit or other high capacity link.

The overall architecture of the ADCS according to the present invention is shown in somewhat greater detail and in block diagram form in FIG. 2. At premises 32 of subscriber 30, for example, an antenna 131 is installed for receiving a broadcast signal from transmission subsystem 20. The received signal is routed to an intelligent control box ("ICB") 130, which is also configured to receive input from all other available sources 100. ICB 130 is equipped to receive input from subscriber 30 and to transmit an appropriate signal, such as an NTSC, PAL, SECAM or other analog standard television signal, to television 34. ICB 130 may also be equipped to route signals to one or more additional televisions or other terminal devices, such as personal computers ("PCs") (not shown). ICB 130 may be coupled to a network interface unit ("NIU") 36 in order to communicate over a network. Alternatively, the functionality of NIU 36 could as easily be implemented by a device resident within ICB 130.

Also shown in FIG. 2 as connecting subscriber premises 32 with return path facility 50 by way of the PSTN 40 and telephone service provider facilities 45 is an optional path 42 via which return path facility 50 can send messages to ICB 130. Return path facility 50 receives via telephone service provider 45 messages not only from subscriber 30, as shown, but also from any number of subscribers (not shown in this view) reached by transmission subsystem 20, preferably by a high capacity link of sufficient bandwidth to simultaneously accommodate messages from a large number of subscribers.

I. Program subsystem

As shown schematically in FIGS. 1 and 2 and functionally in FIG. 3, and as shown in still greater detail in FIG. 4, program subsystem 10 is equipped with conventional equipment to receive content via any suitable communications link from all programming and data sources. Programming and data from non-local content providers 12 can be received, for instance, via satellite (see FIG. 1) by tuner 23A, while programming from local TV content providers 14 could be received by tuner 23B from other transmission media, such as microwave or cable, and also by video tape, compact disc ("CD") delivered to the program subsystem. Among the variety of sources of content to which program subsystem 10 is preferably linked are video and audio entertainment programming, data services such as electronic catalogs, stock market quotations, sports scores and the like, electronic newspapers and other services, as well as other types of content that are as yet unavailable.

In addition to equipment for receiving the variety of content via various transmission paths, program subsystem 10 includes conventional video server means 28 for providing VoD and/or NVoD services, for example. With VoD and NVoD, a preselected set of films is delivered simultaneously, but out of phase (i.e., staggered) by a predetermined amount of time, in order to permit VCR-like functionality to subscribers 30, as discussed in greater detail below. For example, a selection of the 10 top films at any given time may be provided.

Program subsystem 10 tuners 23A and 23B and video server system 28 are coupled to program subsystem bus 21, on which the received data may be read and managed by one or more processors 24. Processor 24 can be implemented by a conventional device or set of devices having sufficient processing power to manage the receipt and suitable signal processing of content, as described, and also to manage the video server system 28. In carrying out its tasks, processor 24, by way of program subsystem bus 21, can read system software and other data from static memory 25, and can

INTRADISTRICT ASSIGNMENT

8. Because this action concerns claims for patent infringement, this case is not subject to assignment to a particular location or division of the Court pursuant to Local Rule 3-2(c).

COUNT ONE: INFRINGEMENT OF U.S. PATENT NO. 5,809,492

9. AT&T incorporates by reference Paragraphs 1 through 8, as if fully set forth herein.

10. United States Patent No. 5,809,492, entitled "Apparatus and Method for Defining Rules for Personal Agents" (hereinafter, the "492 patent"), duly and legally issued on September 15, 1998 after a full and fair examination. AT&T Intellectual Property II, L.P. is the assignee of all rights, title, and interest in the '492 patent, including the right to sue and recover for all past infringement. A true copy of the '492 patent is attached as Exhibit A.

11. As an example, and not to be limited to only a single infringing product, TiVo has developed, has tested, promotes, markets, and sells DVR units that can execute software program modules for managing the recording and deletion of programs, such as "Season Pass Manager." Such products (or the use or operation of such products) fall within the scope of one or more claims of the '492 patent.

12. TiVo has infringed and continues to infringe the '492 patent, by, among other acts, making, using, offering for sale, selling, and/or importing within this Judicial District and elsewhere in the United States, without license or authority by AT&T, products covered by one or more claims of the '492 patent, including, but not limited to, DVR units that can execute software program modules for managing the recording and deletion of programs.

13. As a consequence of TiVo's infringement, AT&T is entitled to recover damages adequate to compensate it for the infringement complained of herein, but in no event less than a reasonable royalty.

14. TiVo has caused and will continue to cause AT&T substantial damage and irreparable injury by virtue of its past and continuing infringement of the '492 patent. TiVo will

store and retrieve data in dynamic memory 26 and mass storage device or devices 27, all of which can be implemented with conventional technology.

Additional input to program subsystem 10 originates at return path facility 50 and is provided over link 52, which can be a broadband link or other link suitable for carrying a large volume of data. The data provided to program subsystem 10 by return path facility 50 includes, but is not limited to, requests for authorization for requested programming, such as pay-per-view, VoD, or NVoD. According to known techniques, the processor or processors of program subsystem 10 includes functionality for receiving requests for pre-stored authorization codes corresponding to the subscribers and for incorporating one or more codes into the broadcast signal in order to enable a requesting subscriber to receive an encrypted program, as will be further described below. Input received over link 52, in the form of authorization requests or messages, is demodulated or otherwise processed as necessary by I/O device 22A, and is provided on program subsystem bus 21 for processing, management and storage by processor 24. Messages can be sent by program subsystem 10 via I/O device 22A back to return path facility 50 as necessary, such as to acknowledge an authorization request or to provide information for use in managing the ADCS.

Referring again to FIG. 2, the primary functionality of the program subsystem 10 is shown. Incoming channel service, VoD or NVoD sources, and all other input content are received according to known methods. The received programming and content are then processed to produce a signal that can be provided via link 18 to transmitter subsystem 20 in a form suitable for physical transmission. The processing performed by program subsystem 10 applies known methods to achieve a maximum transmission rate and a minimal probability of transmission errors, while keeping to a minimum the amount of transmitter power required, and the RF bandwidth required.

The available bandwidth for transmission is likely to be driven in part by the regulatory environment, as well as by the availability (or lack of availability) of suitable UHF channels in a given market. Under the present regulatory framework, channels are defined in 6 MHz RF bands. The UHF spectrum, defined as the 470-806 MHz band, is divided into 56 such 6 MHz channels, which are identified by convention as numbers 14 to 69. The UHF band has been used to provide television service for decades, thus transmission equipment is available from a variety of vendors.

Propagation of signals in the UHF spectrum is generally line-of-sight, and thus limited by obstructions in the path between the transmitter and a receiver, such as mountains and the curvature of the earth. However, the presence of these obstructions can lead to the diffraction of UHF signals and thus to a certain degree of circumvention of the obstacles. Signals in the UHF spectrum also reflect off certain obstructions, which tends to divide the energy in the signals into fractions that propagate to a given receiver over more than one path. Such multipath propagation can lead to fractions of a UHF signal arriving at a receiver at slightly different times, causing a phenomenon known as "ghosting," in which one or more phantom images trail a primary image on the television screen. Ghosting can be addressed through channel equalization techniques.

Transmission of a large number of component signals in a single-to-multipoint UHF transmission poses a number of technical and regulatory problems. Present FCC allocation rules, for example, grant RF spectrum rights to one use or user exclusive of others and in such a manner as to be protected from interference. Co-channel interference is addressed by enforcing a mileage separation of broadcast facilities to hold the desired to undesired signal ratio (D/U) to a predetermined level.

Digital transmission is less susceptible to co-channel interference. Nevertheless, it is expected that existing NTSC (or PAL, SECAM or other analog television standard) channels will be protected to the extent presently required for a number of years.

In addition to co-channel protection, UHF signals are protected by rules designed to limit interference between adjacent channels. These rules, known as "UHF taboos," constrain the D/U between immediately adjacent channel signals to be -6 dB or greater, require separation of transmitter sites by a predetermined distance to address interference introduced by local oscillator radiation, require separation of image frequencies associated with the visual and sound carriers, and preclude use of the second through fifth adjacent channels to combat intermodulation distortion. The UHF taboos are set forth in 47 C.F.R. § 73.699 (Table II)(incorporated herein by reference).

In order to support the delivery of the maximum number of programming channels, as well as a desirable number of films available for VoD or NVoD delivery, and permitting this service to be as close as possible to "VCR functionality," it is desirable to substantially reduce the bandwidth necessary for the delivery of adequate service. Several approaches may be used.

In one embodiment of the system and method according to the present invention, the resolution of the programming is purposely selected to be equivalent to that available from video home system (VHS) video tapes or other publicly acceptable medium having resolution lower than is associated with conventionally broadcast television signals. The resolution of VHS-quality video is approximately 256 by 240 pixels. Since video consumers appear comfortable with this resolution, transmission by the ADCS program and transmission subsystems would appear to be efficient yet unobjectionable at that degree of resolution.

In another embodiment, a data compression scheme is used, for example, MPEG2 compression (Moving Pictures Expert Group standard 2, an international video and audio compression and transmission standard described in ISO/IEC CD 13818-1, the contents of which are incorporated by reference herein) or any other suitable standard, format, protocol, data structure, sequence or organization scheme for reducing the bandwidth required for transmission. Assuming video stream(s) encoded at 3 Mbps, the number of such digital video streams available in a 6 MHz channel can be computed according to the following relation:

$$R = B \cdot \log_2(N) \cdot U$$

where
 R=bit rate;
 B=channel bandwidth;
 N=signalling (or quantizing) level of modulation;
 U=payload useability (to take into account forward error correction);

Quantizing Level N	$\log_2(N)$	3/4			Bit Rate	Number of Video Streams
		Payload	Bandwidth	Bit Rate		
2	1	0.75	6	4.5	1	
4	2	0.75	6	9	3	
8	3	0.75	6	13.5	4	
16	4	0.75	6	18	6	
32	5	0.75	6	22.5	7	
64	6	0.75	6	27	9	
128	7	0.75	6	31.5	10	
256	8	0.75	6	36	12	

The foregoing relation applies to any video encoding/compression scheme and the results in the table apply for any video streams encoded at 3 Mbps per video stream.

The quantizing level and the efficiency with which the video streams are encoded are functions of available technology. In general, programming that will be provided as channel service is digitized, suitably compressed, encoded and modulated according to known techniques, some alternatives for which are shown in FIG. 5. To begin with, the multiple video sources received by program subsystem 10 may each be source coded 112 (or "compressed") if possible in order to reduce the bandwidth needed for transmitting them. Suitable source coding techniques that may be used include character coding, sampling, quantization, pulse code modulation (PCM), differential PCM (DPCM), block coding, and synthesis/analysis coding. Another class of source coding or compression techniques, known as redundancy reducing coding, reduces the volume of data required to transmit a signal by eliminating redundancy in the signal, as well as information that would be rejected rather than processed by a viewer because of psychobiological limitations on human perception. For example, a typical scene in a video contains much information that does not change from frame to frame. Transmitting only the changes in the scene permits a significant reduction in the volume of data transmission without any loss of information to the viewer. Another form of source coding or compression takes advantage of the psychobiological phenomenon that the human eye is less capable of resolving colored images than black and white images, permitting a commensurate reduction in the amount of data to be transmitted without a detectable loss of resolution.

Well-known compression techniques may be used. For example, but without limitation, the MPEG-1 (Motion Picture Experts Group of the International Standards organization (ISO)), or MPEG-2 standard would offer the possibility of "scalable" resolution, and which is currently being implemented in the form of a commercially available integrated circuit chip. However, any other suitable compression technique could also be used. Using MPEG-1 merely as an example, a compression ratio on the order of 100:1 can be achieved. Accounting for imperfections in transmission, the effective bandwidth required for a video stream can accordingly be reduced to the order of 0.08 MHz, implying a theoretical upper limit of 72 VHS-quality video streams per 6 MHz channel according to the current spectrum allocation.

Suitable transmission efficiencies may be achieved with currently available channel coding methods, such as 45 16-QAM (quadrature amplitude modulation), 4-VSB (vestigial sideband), 1-PSK (phase-shift keying) or OFDM (orthogonal frequency division multiplexing). The second processing step for the incoming video streams is channel coding 114, which can be used to reduce required bandwidth and the presence of transmission errors. Channel coding includes waveform coding, such as M-ary signalling, and antipodal, orthogonal, biorthogonal and transorthogonal coding. Channel coding also can include structured sequence coding, for example using block codes or convolutional codes. Both block codes and convolutional codes are directed to minimizing the bit error ratio (BER) via forward error correction (FEC). BER is one of the most important quality factors observed in evaluating digital transmission systems. FEC techniques are intended to reduce residual BER by several orders of magnitude and also to increase system gain by encoding the bit stream prior to modulation. The coding involves adding extra bits to the bit stream according to specific rules; thus, they intentionally add a certain amount of redundancy. Using the foregoing techniques, reliability can be traded off against efficiency to achieve suitable system performance.

After the incoming signal has been suitably source coded and channel coded, the compressed multiple video sources may be further processed using spreading or spread spectrum techniques 116, such as direct sequencing (DS), frequency hopping (FH), time hopping (TH) and hybrids of these techniques. Spread spectrum techniques improve a signal's interference resistance and thus bandwidth efficiency characteristics by distributing the transmitted power over a bandwidth sufficiently wide to ensure that the power per unit bandwidth is kept very small.

Spread spectrum signals are difficult for a casual listener to intercept, but true security requires encryption. In order to protect them against unintended reception, the multiple video source signals are encrypted 118 according to known methods, such as block and data stream encryption techniques.

The digital multiple video source signals, having been compressed, channel coded and spread, are also synchronized 120 and modulated 122 according to known methods, some of which methods are listed, without limitation, in FIG. 5.

Finally, the compressed, channel coded, spread, encrypted, synchronized and modulated multiple video source signal is multiplexed 124 according to known techniques listed at reference numeral 124 of FIG. 5, including, without limitation, frequency division multiplexing, time division multiplexing, code division multiplexing, space division multiplexing and polarization division multiplexing. The fully processed, modulated and multiplexed signal is provided, as shown in FIGS. 1 and 2, over link 18 to transmitting site 20.

The techniques referred to in the schematic of FIG. 5 and the accompanying description, are illustrative but not exhaustive of the techniques that could be used to practice the present invention. Other known techniques may also be used, and their choice, as well as chosen parameters, may depend not only on designer preferences but on design constraints imposed by the setting of the transmitter, the chosen equipment, and other factors. Moreover, the processing associated with these techniques could be performed in program subsystem 10, as described above, but could alternatively be done, at least in part, at transmission subsystem 20 as shown in FIG. 6 and described in the text accompanying that figure.

The VoD or NVoD functionality of the program subsystem may also be implemented according to known methods. VoD service is widely understood to mean that a desired video program can be viewed within 5 minutes after it has been selected. Preferably, a user can exercise virtual control over the transmission by pausing, rewinding, fast forwarding or other function as one would do with a conventional video cassette player. In order to achieve this functionality, for example with a two-hour long film, it would be necessary for twenty four simultaneous video streams to be transmitted, one beginning anew every five minutes. Fast forwarding, rewinding and pausing are thus achieved by tuning to the appropriate time-shifted channel. The program subsystem 10 would receive a request for authorization from a subscriber 30, as further described in connection with FIG. 9, and would transmit this authorization, for example in a vertical blanking interval of a preselected channel. The authorization would be received by the ICB 130 of the requesting subscriber 30 and thereby enable receipt and viewing of the appropriately time-shifted channel. Providing 10 film offerings of two-hour duration with 5 minute VoD functionality would require 240 simultaneous video streams. Provided that 8 video streams can be reliably transmitted in

a single 6 MHz RF channel, 30 RF channels will be required for this VoD function, the availability of which channels will depend on the market, as well as the state of applicable federal regulations. To facilitate the provision of VoD, relaxing the delay from 5 to 10 minutes, thus being closer to NVoD, would reduce by half the required total bandwidth to 15 RF channels.

II. Transmission Subsystem

Transmission subsystem 20 is shown in detail in FIG. 6. Although the multiplexing and modulation of signals has been described above in connection with program subsystem 10, that functionality, or some of it, could alternatively be provided at transmission subsystem 20, as shown in FIG. 6. If the functionality were provided in the program subsystem 10, it could accordingly be excluded from transmission subsystem 20. As shown herein, however, transmission line 18, which may include a plurality of separate lines, delivers the entirety of the programming content from program subsystem 10 to transmission subsystem 20. Subsets of the programming content may be input to multiplexers 21A, 21B, 21C, . . . , 21N, where they are multiplexed according to known methods, as described in connection with box 124 of FIG. 5. The multiplexed signals output from multiplexers 21A, 21B, 21C, . . . , 21N are then digitally modulated at 22A, 22B, 22C, . . . , 22N, respectively, according to conventional modulation techniques. Each multiplexed, modulated signal is then upconverted to RF frequency by a corresponding upconverter 23A, 23B, . . . , 23N, which may be implemented by conventional equipment. Output from upconverters 23A-23N is then amplified by power amplifiers 24A, 24B, 24C, . . . , 24N, respectively. Each power amplifier 24A-24N, which may be conventional equipment, should be capable of generating a peak output sufficient for appropriate radiated RF energy to allow proper reception of equivalent NTSC (or PAL, SECAM or other analog television standard) transmission. Output from each power amplifier 24A-24N is filtered by one or more filters 25A, 25B, 25C, . . . , 25N, each filter being of a conventionally available sort and sufficiently sized to accommodate the power of the output of the corresponding amplifier 24A-24N, to reduce spurious signals to acceptable levels. In an alternative embodiment, the number of sets of transmitting equipment could be reduced by further multiplexing a set of signals prior to modulation, upconversion and amplification.

The output signals from the filters 25A-25N are applied to combiners, which may be conventionally available equipment. For example, output from filters 25A and 25B are combined by combiner 26AB, while output from filter 25C and perhaps other filters are fed to combiner 26C, and output from filters including filter 25N are input into combiner 26N. The output of each combiner 26AB, 26C, . . . , 26N is applied, if necessary, to corresponding transmission lines 27AB, 27C, . . . , 27N, respectively. Since the transmission lines may be required to transmit several high-powered signals, conventional coaxial transmission lines may not be useable. In such cases, transmission lines are preferably waveguides of large diameter (e.g., 18 inches in diameter). Waveguide size is driven by the lowest frequency to be carried, while the highest frequency that can be carried by the same waveguide is limited by the transmission efficiency, which may be about 10% above the lowest channel frequency. Since the UHF television band straddles 470-806 MHz, as many as five different waveguides would be needed to support transmission across that band, although for convenient illustration, three are shown.

Portions of waveguide transmission lines 27AB, 27C and 27N are supported by at least one transmission tower 28

having a suitable height above the ground, while taking into account the relative altitude of its base. For example, transmission tower 28 could be 1000 feet high or more, if required. Each of the waveguide transmission lines 27AB, 27C and 27N is coupled to a corresponding antenna 29AB, 29C and 29N, respectively. These antennas are also supported by transmission tower 28, which must be sufficiently strong to bear them along with the waveguide transmission lines 27AB, 27C and 27N, and should accordingly be fabricated with sufficient strength according to known methods. Preferably, at least three antennas should be used in order to support transmission across the entire UHF band. As with transmission lines, antenna efficiencies are frequency-dependent. The height of transmission tower 28 should ensure sufficient vertical aperture of antennas 29AB, 29C and 29N. As an alternative to a large single structure, several smaller supporting structures could be used (not shown). The transmission subsystem 20 can be expected to generate a high level of non-ionizing radiation, which, depending upon the height at which the antenna(s) are actually mounted, may require that its site be located most preferably at an appreciable physical distance from human populations.

As described above, known components can be used in transmission subsystem 20, preferably of a type suitable at least for use with high definition television (HDTV).

III. Return Path Subsystem

Broadcasts by the transmission subsystem 20 are received by subscribers who are within the broadcast pattern of the transmission subsystem 20 and who at their premises 30 have a proper receiving device. According to the present invention, the receiving device is a component of a piece of equipment having a number of functions, and which is referred to as an ICB (intelligent control box) 130.

Along with the return path facility 50 to which it is linked by a telecommunications system, such as PSTN lines 40 (or a suitable wireless alternative, not shown) and switching and other facilities operated and maintained by public telephone services provider 45, the ICB 130 constitutes a portion of the return path subsystem 80 of the described embodiment of the system according to the present invention. The use of the PSTN lines 40 between the portions of the return path subsystem 80 takes advantage of the realization that the timing and information content of communications by humans differ enormously from those of the broadcast path, perhaps on the order of 10^9 . Information sent upstream by a subscriber 30, moreover, would tend to be sporadic rather than continuous. The second portion of the return path, return path facility 50, implements two types of functionality: switching and accounting system 60 and external transmission system 70, each of which will be described below.

A. Intelligent Control Box (ICB)

ICB 130 forms that portion of the return path subsystem 80 of the system according to the present invention that resides at the premises of subscriber 30 and receives the digital UHF broadcast from the transmission subsystem 20. In addition to this role, ICB 130 also provides subscribers 30 with a variety of additional capabilities and functions.

The architecture for an embodiment of ICB 130 according to the present invention is shown in FIG. 7. ICB 130 acts as a matrix switch (i.e., an electronic gatekeeper) to provide an interface between the ADCS, the full variety of other programming sources, and the viewer's television monitor, computer, or other peripheral devices. In addition to over-the-air UHF ADCS broadcast, ICB 130 is provided with input ports to receive signals from any number of available sources. ICB 130 can accept input from an antenna 132, which can carry NTSC signals as well as digitized UHF

signals transmitted by transmission subsystem 20. According to the present invention, ICB 130 can also accept input from any number of non-ADCS inbound program sources, such as from a VCR 133, from a DBS provider 134, from an MMDS provider 135, from an LMDS provider 136, or from a CATV provider 137. Input from any other available source, including a mass storage or other device, as well as any presently existing or future transmission type, may be received by ICB 130, as denoted by the input port identified as "other" 138, which may be adapted or retrofitted to receive an appropriate connector as necessary.

All program inputs received by ICB 130 are transmitted to input selector (matrix switch) 170, which can be any available hardware for switching between a plurality of signals, most preferably under the control of a conventional or custom processor. Input selector 170 is coupled to system control bus 180 of ICB 130 by way of which it can receive switching signals from CPU 184 in order to execute the selection commands of subscriber 30, as further discussed below.

ICB 130 may include a tuner for each input type. Signals switched by input selector 170 are each provided to an appropriate tuner. A selected over-air-air NTSC signal received via antenna 131, for example, can be tuned by analog tuner 161, as can a selected CATV signal received at port 137. In an alternative embodiment, one or more additional analog tuners could be provided for simultaneously receiving two analog signals, if it is desired to route two such signals to different televisions, computers or other devices, to provide a picture-in-picture (PIP) display, or to support other functionality.

Signals on input 132, such as the digital UHF signal transmitted by transmission subsystem 20 according to the present invention are received by ATV tuner 162, which can be a conventionally available digital tuner and demodulator. DBS input arriving at port 134, if selected and accordingly switched by input selector 170, can be received by tuner 163, while MMDS signals entering at port 135 and switched by input selector 170 are received by MMDS tuner 164. Similarly, LMDS signals arriving at 136 and switched by input selector 170 can be tuned by LMDS tuner 165. All of the foregoing tuners may be implemented by any commercially available devices for receiving the respective signals. Moreover, the foregoing tuners are illustrative rather than exhaustive; signals requiring tuning that cannot be done by way of the foregoing tuning devices, such as signals according to as yet unconceived schemes, can be received by a suitable associated tuning device or devices represented by box 166 with which ICB 130 can be retrofitted as desired. Subscriber 30 thus does not have to collect, operate and maintain a set of devices from various vendors each of which provides only a single purpose solution.

The output from analog tuner 161 (and any additional analog tuners that may optionally be included in ICB 130) may be transmitted to secure processor subsystem 172. Secure processor subsystem 172 may be any known system for receiving scrambled or otherwise secured signals over a private distribution system. Typically, the secured processor subsystem 172, its components, and their functionality are compatible with encryption techniques of the sort employed in program subsystem 10, illustrated, for example, at block 118 of FIG. 5. Secure processor subsystem 172 includes a conventional analog descrambler 173 for descrambling signals received and forwarded by analog tuner 161. Secure processor subsystem 172 (or other processor means described below) is programmed to detect one or more authorization codes included in received transmissions and

only to descramble a particular transmission on detection of such code or codes. Secure processor subsystem 172 may include a non-volatile memory subsystem 174 for storing information necessary for decoding received transmissions, including normal cable program transmissions, pay-per-view program transmissions and the like.

In order to interface televisions or other viewing devices accepting a baseband input with a decoding device (for example, a CATV decoder, not shown), secure processor subsystem 174 is optionally coupled to EIA-563 standard baseband interface, the baseband output of which it can provide to baseband-equipped televisions or other devices over line 182B.

Turning to the digital signals received by ICB 130, each tuner among the set of tuners 162-166, demodulates signals it is tuned to receive in order to reconstruct the desired analog signal. The resulting signal from each tuner is transmitted by a dedicated line (here shown in a bundle of such lines denoted in FIG. 7 using the notation "/5") to video

decoder 176, a conventionally available device for interpreting an input stream according to a preselected technique compatible with the method used for compressing the video signal at the program subsystem 10, for example. Video decoder 176 thereby recovers the compressed signal to produce a video signal. An analog channel received for example by analog tuner 161 and selected by input selector 170 may be routed via secure subprocessor system 172 directly to NTSC signal generation module 182. Digital signals received by tuners 162-166 may also be provided on dedicated lines (shown in bundled form using the notation "/5") to system control bus 180, via which it could be provided, as one illustrative example, via computer interface 196 to a computer or other device.

Analog descrambler 172 and video decoder 176 are both coupled to graphics assist module 178. Graphics assist module 178 can be any conventional graphics processor suitable for adding a desired graphics overlay onto the video signals received from the video decoder 176 and/or the analog descrambler 172 in response to instructions implemented by software running on CPU 184 and provided over system control bus 180. The functionality of graphics assist module 178 could alternatively be implemented by CPU 184, if it were suitably programmed and capable of sufficient throughput. As will be further described below, graphics assist module 178 generates a video signal overlaid on the existing television signal providing, in response to instructions received over system bus 180. For example, a menu of viewing options, information regarding the ordering of a VoD or NVoD video, or any other information regarding subscriber transactions may be overlaid on the screen. Graphics assist module 178 may generate a menu, a button image, or a logo or other pattern or symbol, and inject the image into the video stream such that the image appears in a designated portion of the screen. Information used in generating the screen images, as described below, may be available in received signals or may use information stored in the ICB 130, for example, pre-loaded in static memory 186.

The video stream emerging from analog descrambler 173 or video decoder 176, and from graphics assist module 178, is provided to NTSC signal generation module 182. As an alternative to NTSC, signal generation module 182 could generate signals according to any available and desired video format. NTSC signal generation module 182 includes conventional circuitry for transforming the output of analog descrambler 172, video decoder 176 and graphics assist module 178 into a standard form usable by a conventional

15

television or monitor (not shown in this view). This functionality could, alternatively, convert input signals to any suitable standard format. As shown in FIG. 7, NTSC signal generation module 182 generates a set of coordinated signals and provides them to one or more terminal devices (not shown) adapted to receive and display television signals, including without limitation one or more televisions or personal computers. Output from NTSC signal generation module 182 to the one or more terminal devices includes an NTSC video stream 182A on channel 3 and/or 4, a composite/Svideo Out signal 182B (which may also carry baseband signals routed through EIA-563 171), and left and right audio signals 182C. The signal provided by ICB 130 could be any signal deriving in whole or in part from received input and be in any desired form. Even more generally, CPU 184 can issue instructions to manage televisions, computers, or other terminal or peripheral devices without limitation, for example, via computer interface 196, which can be any suitable conventional device, coupled to system control bus 180.

In addition to the previously described components and functionality, which have to do with the receipt and provision of television signals, possibly with inserted graphics, ICB 130 includes a number of additional components for implementing other functions. The processing associated with managing and controlling the functions of the ICB 130 is performed by CPU 184. CPU 184 may be any suitable commercially available processor, and preferably one capable of performing in excess of 30 Mips, such as a POWER PC® or PENTIUM® integrated circuit (IC) chip, programmed to process the incoming digitized UHF signal according to known methods, as well as selecting and processing the remainder of the input provided at ports 134, 138, 140, 144 and 145, when chosen by subscriber 30. Alternatively, CPU 184 could be a custom chip or chip set for carrying out the same set of functions. In running system software stored in static memory 186 (which may be an EEPROM, ROM or the like), CPU 184 has access via system control bus 180 to static memory 186 and to dynamic memory 188, which may be conventionally available memory, preferably with a storage capacity of at least about 4-5 MB. CPU 184 also has access via system control bus 180 to information regarding the current channel, to the graphics assist module 178 regarding graphical display signals, and to video decoder 176 and input selector 170, as well as to secure processing subsystem 172, and the tuners 161-166, in order to assert appropriate command messages to those components.

Interaction by subscriber 30 with ICB 130 may be by any suitable means, but is preferably by conventional infrared (IR) remote control, as in the described embodiment. Input controller 190 receives signals from subscriber-controlled remote control input device 158 via an IR receiver and associated circuitry (not shown). Using input device 158 (or any other suitable input means), subscriber 30 may issue to the ICB 130 instructions to switch between programs, to request VoD or NVoD services via the return path facility 50, to control the delivery of such requested services with VCR functionality, to purchase items offered for sale on shopping channels, or to provide any desired input originating with subscriber 30. Instructions received by input controller 190 may be asserted via system control bus 180 as commands to input selector 170 and tuners 161-166, to secure processor subsystem 172, as well as to video decoder 176, and graphics assist module 178. Input controller 190 can also send messages to CPU 184 when, for example, subscriber 30 requests an interactive program or other service.

16

Interaction by the ICB 130 with return path facility 50 via PSTN lines 40 and telephone service provider 45 is conducted by CPU 184 through WAN communications processor and modem 192. WAN communications processor and modem 192 can be any conventional device suitable for converting messages received via system control bus 180 into a WAN-compatible protocol, such as IP (internet protocol), and modulating the converted message signals received from WAN communications processor for transmission over the PSTN lines 40 and via switching and other facilities maintained and operated by a telephone service provider 45 (as shown in FIGS. 1-3).

In addition to the primary function of delivering television programming to subscriber 30, or management instructions to one or more terminal devices such as televisions or computing devices, CPU 184 operating according to software residing in static memory 186, for example, and using known methods, implements an authorization function enabling subscribers (and only subscribers) of particular services, to be able to view programming associated with those services. CPU 184 is also programmed to capture and send via telecommunication connection 194 to return path facility 50 subscriber requests for transactions input via remote control system 158, as well as data regarding viewership and other observable events.

B. Return Path Facility

A second portion of the return path subsystem 80 is a return path facility 50. Return path facility 50 is shown in FIGS. 1-3 in schematic form, illustrating the functionality of that component of the system according to the present invention. Referring to FIG. 8, return path facility is shown in greater detail.

In brief, return path facility 50 provides an external transaction system for facilitating viewing requests, purchases and other transactions by subscriber 30 executed from the premises via the ADCS according to the present invention, system, and also for facilitating fulfillment of the viewing, purchase, or other request, including taking an order, processing the order, establishing delivery of the item requested or purchased, conducting billing for the transaction or purchase, and sharing of the proceeds with the seller. In addition, return path facility 50 also may capture viewing and purchase pattern information for subsequent compilation by a market analysis center to determine types of content most valuable to each individual customer. This process facilitates the transmission of targeted advertising messages that can trigger interactions for more information or a transaction.

The return path subsystem 50 is ultimately adaptable from an ADCS to a symmetrical data communications system when the PSTN or other non-broadcast return path has sufficient available bandwidth to accommodate full video delivery. Return path facility 50 includes conventional processing means 54, which can be one or more commercially available processors and which is coupled via a return path facility bus 51 to at least one static memory device 55, at least one dynamic memory device 56, and which may also be coupled to at least one mass storage device 57. These memory devices can all be implemented with commercially available hardware.

Two way communications between return path facility 50 and subscribers, such as illustrative subscriber 30, can be maintained over a public telephone or other transmission line or lines via I/O device 53, which can be any conventional device suitable for modulating and demodulating input and output in a manner consistent with the transmission line. Since the return path subsystem 50 is responsible

for receiving request messages, demographics and other information from all of the subscribers within the signal range of transmission subsystem 20, the transmission line or lines are preferably of capacity sufficient to handle this traffic.

According to functionality described at greater length below, return path subsystem 50 sends messages, such as authorization requests, to program subsystem 10 via I/O device 58 on return path subsystem bus 51. Optionally, messages such as demographic, billing or other transaction data can also be sent to other service providers 48 via I/O device 59 on return path subsystem bus 51. I/O devices 58 and 59 can be implemented by any conventional equipment suitable for this purpose.

As shown in FIG. 3, the functionality of return path subsystem 50 includes two main modules: session control module 60 and administrative module 70. Processing associated with both modules is performed by processor 54 according to software that may be pre-stored in static memory 55, and utilizing dynamic memory 56 and mass storage 57 as necessary. In the illustrated embodiment, but without limitation, session control module 60 of return path facility 50 is responsible for maintaining an interface with each ICB 130 that has initiated a session. Requests for programming or other transactions are managed and monitored, as are audience observations provided by ICB 130 according to a desired, predetermined approach. Session control module 60 also includes the recording of desired aspects of messages and transactions in mass storage device 57. Also, transaction requests, demographic or other marketing data, or other received messages are forwarded for processing by administrative module 70.

Administrative module 70 captures the occurrence of requested transactions for purposes of billing, recording viewing patterns, and for fulfilling the requested transaction by transmitting appropriate requests as necessary over link 52 to program subsystem 10. Requests for other sorts of transactions, such as purchasing offerings on home shopping channels or any other transaction may be accounted for and then provided to the appropriate one of a set of other service providers 48 with which the operator of the ADCS has established a relationship. Administrative module 70 also gathers customer information regarding customer viewership, purchasing patterns, and demographic or psychological information important for marketing and stores this information as necessary and/or provides it to one or more of the other service providers 48 according to a predetermined arrangement. Such information gathering may be done, for example, but without limitation, pursuant to agreements with subscribers permitting the gathering of such information, possibly with a reduction in subscription rate or some other consideration to the subscriber. Administrative module 70 of return path subsystem 50 according to a predetermined scheme can also keep track of viewing and transaction requests for one or more other service providers 48, and can use this information to conduct billing on behalf of those providers 48.

IV. Operation of the ADCS System According to the Present Invention

A logic flow diagram for processing associated with a representative transaction performed by an embodiment of an ADCS according to the present invention is shown in FIG. 9. In particular, subscriber 30 via ICB 130 initiates a request for a particular program requiring authorization, and the ensuing processing or activity by the ICB 130, the return path subsystem 50, program subsystem 10, and transmission subsystem 20 are described. On the top edge of each box is

indicated the subsystem of the ADCS responsible for performing the indicated function.

Process 200 begins at step 202, at which the ICB 130 tuner (i.e., the input selector 170) is tuned to regular programming selected by subscriber 30, which may be a default setting that could be either factory set or determined by subscriber 30. Using the IR remote 158, the user may request a viewing session, at 204. CPU 184 of ICB 130 retrieves from static memory 186 (or other storage not shown in FIG. 5) pre-stored code for a viewing selection list and performs any necessary processing. Then, via system control bus 180 and, if necessary, video decoder 176 and graphics assist 178, the pre-stored selection list code is processed by NTSC signal generation module 182 and provided over line 182A to a terminal device, such as television set 34, at step 208.

If subscriber 30 expresses no interest in a transaction, process 200 simply returns to the beginning, step 202. If subscriber 30 requests a transaction however at 210 (using IR remote 158) and IR controller 190 receives the request, at 212, the CPU 184 of ICB 130 retrieves from static memory 186 a pre-stored request for a personal identification number (PIN), which it then causes to be displayed on television set 34 as, for example, the character string "Enter PIN." The subscriber 30, having been prompted by the message, via IR remote 158 enters his or her PIN, which is received by IR controller 190 and stored by CPU 184 in dynamic memory 188. The CPU 184 of ICB 130 according to instructions it retrieves from static memory 186 initiates a procedure to verify the PIN, at 224.

If the PIN is not verified, CPU 184 causes to be displayed on television set 34 the legend "Wrong PIN, Cancel?", or other legend to the same effect, and waits for a response from subscriber 30. If subscriber 30 responds in the negative, process 200 returns to step 214, in which the PIN is again requested (a memory or input error presumably having been made by subscriber 30). If subscriber 30 responds in the affirmative, process 200 simply returns to initial step 202.

If, on the other hand, the PIN was verified at step 224, then CPU 184 invokes WAN communications processor and modem 192, CPU 184 dials return path subsystem 50 (which might also be referred to as "head end", as it is in FIG. 9). Processor 54 of return path subsystem 50 establishes a session with ICB 130 and reads subscriber 30's request, PIN, and other desired subscriber information. Processor 54 of return path subsystem 50, invoking the administrative module 70, generates billing data associated with this request.

For purposes of this logic flow, "headend" is also, for convenience, intended to include functionality implemented by program subsystem 10. At step 230, the processor 54 of return path subsystem 50 generates and transmits to program subsystem 10 a message that subscriber 30 is authorized for a particular selected viewing choice. An authorization code, which may have been pre-stored at return path facility 50 and retrieved for program subsystem 10, or which could alternatively have been pre-stored in program subsystem 10, is then included in the transmission by program subsystem 10 to transmission subsystem 20 and is thus broadcast, for example in the vertical blanking interval of a particular video stream. The authorization code is inserted into a video stream being transmitted in a frequency band to which ICB 130 is programmed to tune to look for this information. Accordingly, the authorization code is identified in that band by the ICB 130 and extracted according to known methods, at step 232, for use in enabling the decoding of the selected transmission, or tuning to an authorized channel, step 234.

CPU 184 of ICB 130 then ensures that the ATV tuner 162 is tuned to the appropriate band, if necessary, that the input selector 170 is switching the transmission from the ATV tuner 162 through to the video decoder 176, and that the video decoder 176 has the necessary authorization for descrambling the authorized transmission. At 238 the subscriber 30 may end the session using IR remote 158, either at the end of the program, or for any other reason. If subscriber 30 during the course of the transmission does not indicate that he or she wishes to terminate the session, CPU 184 continues to cause the authorized channel to be displayed on television set 34. During this period, other functionality could be represented, including VCR functionality or other subscriber interaction. If, on the other hand, subscriber 30 does end the session, or if it terminates on its own, process 200 in that case returns to initial step 202.

The foregoing describes a preferred embodiment of the present invention. Various changes and modifications to what is disclosed may be adopted or implemented without departing from the scope or spirit of the invention.

What is claimed is:

1. A system for communicating content to a plurality of subscribers, the content carried by signals originating from a plurality of content providers, the system for communicating content comprising:

a plurality of receivers, each receiver for receiving at least one of the content-containing signals transmitted by one of the plurality of content providers;

means for digitizing the signals received from the plurality of content providers, the digitizing means coupled to the receiving means for receiving the signals;

means for compressing the digitized signals, the compression means coupled to the digitizing means for receiving the digitized signals;

a multiplexer coupled to the compression means for multiplexing the digitized signals into at least one digital data stream;

modulating means coupled to the multiplexer for modulating the at least one multiplexed digital data stream;

RF upconverter means coupled to the multiplexing means for channelizing the at least one multiplexed digital data stream into the UHF frequency band;

an amplifier coupled to the RF upconverter means for amplifying the at least one UHF multiplexed digital data stream; and

an antenna coupled to the amplifier for transmitting the at least one amplified UHF multiplexed digital data stream;

whereby an aggregation of programming sources may be provided point-to-multipoint to a plurality of subscribers by way of digital UHF broadcast.

2. The communications system according to claim 1, wherein at least one of the signals from the content providers is transmitted via satellite and the at least one receiver includes means for receiving signals transmitted via satellite.

3. The communications system according to claim 2, wherein a second of the plurality of received signals is provided via cable, and at least one of the plurality of receivers is adapted to receive cable transmissions.

4. The communications system according to claim 3, wherein a third of the plurality of received signals is provided via microwave transmission, and wherein at least one of the plurality of receivers includes a receiver adapted to receive microwave transmissions.

5. The communications system according to claim 4, wherein a fourth of the plurality of received signals is

transmitted optically via optical fiber, and wherein at least one of the plurality of receivers includes a receiver adapted to receive optical transmissions.

6. The communications system according to claim 1, wherein the received signal comprises video data, and wherein the at least one compression means is adapted to compress the digitized signal to a resolution of approximately 256 by 240 pixels.

7. A system for communicating with a plurality of subscribers, the communication involving content carried by signals originating from a plurality of content providers, at least one of the plurality of subscribers having access to a telecommunications system, the system for communication comprising:

a program subsystem, including:

a. a plurality of receivers, each receiver for receiving one of the plurality content-carrying signals;

b. processing means coupled to the plurality of receivers and adapted for:

i. digitizing the signals from the plurality of content providers;

ii. inserting desired information into the digitized signals;

iii. compressing the digitized signals; and

iv. multiplexing the digitized signals into at least one data stream;

c. means for modulating the at least one digital data stream, the modulating means coupled to the processing means;

d. RF upconverter means coupled to the modulating means for channelizing the at least one modulated digital data stream into the UHF frequency band;

e. at least one amplifier coupled to the RF upconverter means for amplifying the at least one modulated digital signal;

f. at least one antenna coupled to the at least one amplifier for transmitting the at least one amplified data stream; and

a return path subsystem coupled to the program subsystem and to the telecommunications system, the return path subsystem including processing means adapted for receiving over the telecommunications system communications from the subscribers and communicating with the program subsystem information to be inserted by the program subsystem into the received, digitized signals;

whereby an aggregation of programming sources may be provided point-to-multipoint to subscribers by way of digital UHF broadcast, subscriber input regarding the broadcast can be received, and program transmissions may be modified in response to the subscriber input.

8. A system for communicating with a plurality of subscribers, the communication involving content carried by signals originating from a plurality of content providers, at least one of the plurality of subscribers having access to a telecommunications system, the system for communication comprising:

a. a program subsystem, including:

i. a plurality of receivers, each receiver for receiving one of the plurality content-carrying signals;

ii. processing means adapted for modifying the received signals, the processing means coupled to the plurality of receivers;

b. a transmission subsystem coupled to the program subsystem and receiving from the program subsystem the received signals, the transmission subsystem including:

- i. means for digitizing the signals from the plurality of content providers;
- ii. means for compressing the digitized signals, the compression means coupled to the digitizing means;
- iii. a multiplexer, for multiplexing the digitized signals into at least one data stream, the multiplexer coupled to the compression means;
- iv. means for modulating the at least one digital data stream into the UHF frequency band, coupled to the multiplexing means;
- v. RF upconverter means coupled to the modulating means for channelizing the at least one modulated multiplexed digital data stream into the UHF frequency band;
- vi. at least one amplifier coupled to the RF upconverter means for amplifying the at least one modulated digital data stream; and
- vii. at least one antenna for transmitting the at least one amplified, channelized, modulated, digital data stream; and
- c. a return path subsystem coupled to the program subsystem and to the telecommunications system, including processing means adapted for receiving over the telecommunications system communications from the subscribers;

whereby an aggregation of programming sources may be provided point-to-multipoint to subscribers by way of digital UHF broadcast and subscriber input regarding the broadcast can be taken into account by the communications system.

9. The communications system according to claim 8, wherein the processing means of the return path subsystem is further adapted to communicate with the program subsystem on receiving a communication from a subscriber, and the processing means of the program subsystem is further adapted to modify the received signals on receiving a communication from the return path subsystem, whereby the transmission may be modified in response to a subscriber request.

10. The communication system according to claim 9, wherein the received signals include a plurality of programs, the subscriber request relates to of the plurality of programs, and the modification of the received signals by the program subsystem in response to subscriber requests is the inclusion of an authorization code in the transmission, the authorization code corresponding to the subscriber request for a particular program.

11. The communications system according to claim 8, wherein the program subsystem further comprises a video server system, the video server system including at least one video server and adapted to play at least one of a preselected set of videos, the return path processor means being further adapted to receive requests for transmission of one of the preselected set of videos and to transmit to the program subsystem a message regarding authorization for transmission to the subscriber of the one preselected video, and the program subsystem processing means being further adapted to include in the transmission an authorization code corresponding to the subscriber and to the selected video.

12. The communications system according to claim 8, wherein the return path subsystem further comprises data storage means, and wherein the return path subsystem processing means is further adapted to create a record corresponding to communications received from subscribers and to store that record in the data storage means.

13. The communications system according to claim 12, wherein the record corresponding to communications received from subscribers contains billing information.

14. The communications system according to claim 12, wherein the record corresponding to communications received from subscribers contains subscriber preference information derived from the subscriber communications.

15. A system for communicating content with a plurality of subscribers, the content originating from a plurality of content providers and transmitted by each content provider via receivable signals, at least one of the plurality of subscribers having access to a telecommunications system, the system for communicating content comprising:

- a. a program subsystem, including:
 - i. a plurality of receivers, each receiver for receiving one of the plurality of signals, each signal from one of the plurality of content providers;
 - ii. processing means adapted for modifying the received signals, the processing means coupled to the plurality of receivers;
- b. a transmission subsystem, including:
 - i. at least one means for digitizing the signals from the plurality of content providers coupled to the receiving means;
 - ii. at least one means for compressing the digitized signals, coupled to the digitizing means;
 - iii. at least one multiplexer, for multiplexing the digitized signals into at least one data stream, coupled to the compression means;
 - iv. means for modulating the at least one digital data stream, coupled to the multiplexing means;
 - v. RF upconverting means coupled to the modulating means for channelizing the at least one modulated multiplexed digital into the UHF band;
 - vi. at least one amplifier coupled to the RF upconverting means for amplifying the at least one modulated, channelized, multiplexed digital signal; and
 - vii. at least one antenna coupled to the at least one amplifier for transmitting the at least one amplified, channelized, modulated, digital signal;
- c. a return path facility coupled to the program subsystem and to the telecommunications system, including processing means adapted for receiving over the telecommunications system communications from the subscribers;
- d. a plurality of intelligent control boxes, each intelligent control box at a subscriber's premises, coupled to the telecommunications system, and including:
 - i. a tuner for receiving transmissions from the transmission subsystem;
 - ii. an input device for accepting subscriber requests; and
 - iii. a processor coupled to the tuner and to the input device and adapted to assert subscriber requests at the return path facility via the telecommunications system;

whereby an aggregation of programming sources may be provided point-to-multipoint by way of digital UHF broadcast, a plurality of subscribers can assert requests in connection with the broadcast, and the return path facility can fulfill the requests.

16. A system for communicating with a plurality of subscribers, the communication involving content carried by signals originating from a plurality of content providers, at least one of the plurality of subscribers having access to a telecommunications system, the system for communication comprising:

- a. a program subsystem, including,
 - i. a plurality of receivers, each receiver for receiving at least one of the plurality of content carrying signals;

23

- ii. processing means adapted for modifying the received signals, the processing means coupled to the plurality of receivers;
- b. a transmission subsystem coupled to the program subsystem and receiving from the program subsystem the received signals, the transmission subsystem including:
 - i. means for digitizing the signals from the plurality of content providers;
 - ii. means for compressing the digitized signals, the compression means coupled to the digitizing means;
 - iii. a multiplexer, for multiplexing the digitized signals into at least one data stream, the multiplexer coupled to the compression means;
 - iv. means for modulating the at least one digital data stream into a radio frequency band, coupled to the multiplexing means;
 - v. RF upconverter means coupled to the modulating means for channelizing the at least one modulated multiplexed digital data stream into the radio frequency band;
 - vi. at least one amplifier coupled to the RF upconverter means for amplifying the at least one modulated digital data stream; and
 - vii. at least one antenna for transmitting the at least one amplified, channelized, modulated, digital data stream; and
- c. a return path subsystem coupled to the program subsystem and to the telecommunications system, including processing means adapted for receiving over the telecommunications system communications from the subscribers;

whereby an aggregation of programming sources may be provided point to multipoint to subscribers by way of digital broadcast and subscriber input regarding the broadcast can be taken into account by the communications system.

17. The communications system according to claim 16, wherein the processing means of the return path subsystem is further adapted to communicate with the program subsystem on receiving a communication from a subscriber, and the processing means of the program subsystem is further adapted to modify the received signals on receiving a communication from the return path subsystem, whereby the transmission may be modified in response to a subscriber request.

18. The communications system according to claim 16, wherein the return path subsystem further comprises data storage means, and wherein the return path subsystem processing means is further adapted to create a record corresponding to communications received from subscribers and to store that record in the data storage means.

19. A device for permitting a subscriber to interact with an asymmetrical data communications system, the device coupled and the asymmetrical data communications system both coupled to a telecommunications system, the device having a plurality of ports each for receiving one of a plurality of signals, comprising:

- a. an input selector coupled to the plurality of ports and adapted for switching signals arriving at the plurality of ports according to a subscriber selection;
- b. a plurality of receivers coupled to the input selector each for receiving a selected signal switched by the selector, the plurality of receivers including:
 - i. a tuner for receiving digital RF transmissions in the UHF band; and

24

- ii. a tuner for receiving analog RF signals; and
- c. a video decoder coupled to the digital tuner;
- d. signal generation means coupled to the video decoder;
- e. an input device for subscriber information input;
- f. a processor coupled to the input selector and the video decoder adapted for reading subscriber input information and actuating the input selector to effect subscriber requests; and
- g. a network interface coupled to the processor and to the telecommunications system for transmitting subscriber messages to the asymmetrical data communications system.

20. The device of claim 19, wherein the plurality of receivers further comprises a DBS tuner coupled to the input selector.

21. The device of claim 19, wherein the plurality of receivers further comprises an MMDS tuner coupled to the input selector.

22. The device of claim 19, wherein the plurality of receivers further comprises an LMDS tuner coupled to the input selector.

23. The device of claim 20, wherein the plurality of receivers further comprises an MMDS tuner coupled to the input selector.

24. The device of claim 20, wherein the plurality of receivers further comprises an LMDS tuner coupled to the input selector.

25. The device of claim 21, wherein the plurality of receivers further comprises an LMDS tuner coupled to the input selector.

26. The device of claim 23, wherein the plurality of receivers further comprises an LMDS tuner coupled to the input selector.

27. The device of claim 19 further comprising a secure processor subsystem coupled to the analog tuner for descrambling analog input signals that require descrambling.

28. The device of claim 27 further comprising an EIA-563 baseband interface coupled to the secure processor subsystem.

29. A method for providing asymmetric data communications services, comprising the steps of:

- a. receiving signals from a plurality of content providers;
- b. digitizing the signals from the plurality of content providers;
- c. multiplexing the signals from the plurality of content providers into at least one data stream;
- d. modulating the at least one multiplexed, digital data stream;
- e. channelizing the at least one modulated, multiplexed digital data stream into at least one digital RF signal;
- f. transmitting the at least one digital RF signal to a plurality of subscribers each having a device for receiving and demodulating the at least one RF signal;
- g. receiving over the PSTN a message from at least one subscriber receiving device; and
- h. retransmitting to other service providers each of the at least one messages from the at least one subscriber receiving device.

30. The method of claim 29 wherein the at least one modulated, multiplexed digital data stream is channelized into at least one digital UHF signal.

31. The method of claim 29 wherein the resolution of the service is approximately 256 by 240 pixels.

1 continue to infringe the '492 patent unless enjoined by this Court. AT&T will suffer further
2 damage and irreparable injury unless and until TiVo is enjoined by this Court from continuing
3 such infringement.

4 **COUNT TWO: INFRINGEMENT OF U.S. PATENT NO. 5,922,045**

5 15. AT&T incorporates by reference Paragraphs 1 through 14, as if fully set forth
6 herein.

7 16. United States Patent No. 5,922,045, entitled "Method and Apparatus for
8 Providing Bookmarks when Listening to Previously Recorded Audio Programs" (hereinafter, the
9 "045 patent"), duly and legally issued on July 13, 1999 after a full and fair examination. AT&T
10 Intellectual Property II, L.P. is the assignee of all rights, title, and interest in the '045 patent,
11 including the right to sue and recover for all past infringement. A true copy of the '045 patent is
12 attached as Exhibit B.

13 17. As an example, and not to be limited to only a single infringing product, TiVo has
14 developed, has tested, promotes, markets, uses, and sells services for displaying streaming
15 television programs and movies. Such services (or the use or operation of such services) fall
16 within the scope of one or more claims of the '045 patent.

17 18. TiVo has infringed and continues to infringe the '045 patent, by, among other
18 acts, making, using, offering for sale, and/or selling within this Judicial District and elsewhere in
19 the United States, without license or authority by AT&T, services covered by one or more claims
20 of the '045 patent, including, but not limited to, services for displaying streaming television
21 programs and movies.

22 19. As a consequence of TiVo's infringement, AT&T is entitled to recover damages
23 adequate to compensate it for the infringement complained of herein, but in no event less than a
24 reasonable royalty.

25 20. TiVo has caused and will continue to cause AT&T substantial damage and
26 irreparable injury by virtue of its past and continuing infringement of the '045 patent. TiVo will
27 continue to infringe the '045 patent unless enjoined by this Court. AT&T will suffer further
28

US 6,983,478 B1

Page 2

U.S. PATENT DOCUMENTS

4,720,873 A	1/1988	Goodman et al.
4,816,904 A	3/1989	McKenna et al.
4,890,322 A	12/1989	Russell, Jr.
4,912,552 A	3/1990	Allison, III et al.
5,010,585 A	4/1991	Garcia
5,038,211 A	8/1991	Hallenbeck
5,046,090 A	9/1991	Walker et al.
5,046,092 A	9/1991	Walker et al.
5,055,924 A	10/1991	Skutta
5,173,900 A	12/1992	Miller et al.
5,191,645 A	3/1993	Carlucci et al.
5,208,665 A	5/1993	McCalley et al.
5,247,347 A	9/1993	Litteral et al.
5,249,044 A	9/1993	Von Kohorn
5,251,324 A *	10/1993	McMullan, Jr. 725/14
5,287,181 A	2/1994	Holman
5,335,277 A	8/1994	Harvey et al.
5,339,315 A	8/1994	Maeda et al.
5,343,240 A	8/1994	Yu
5,357,276 A	10/1994	Banker et al.
5,374,951 A	12/1994	Welsh
5,389,964 A	2/1995	Oberle
5,404,393 A	4/1995	Remillard
5,410,344 A	4/1995	Graves
5,446,490 A	8/1995	Blahut et al.
5,446,919 A	8/1995	Wilkins
5,534,911 A	7/1996	Levitian
5,559,548 A	9/1996	Davis
5,585,838 A	12/1996	Lawler et al.
5,585,865 A	12/1996	Amano
5,600,364 A	2/1997	Hendricks
5,608,448 A	3/1997	Smoral et al.
5,619,247 A *	4/1997	Russo 725/104
5,630,119 A	5/1997	Aristides
5,661,516 A	8/1997	Carles
5,724,521 A	3/1998	Dedrick
5,724,525 A	3/1998	Beyers, II et al.
5,752,159 A *	5/1998	Faust et al. 725/105
5,754,383 A	5/1998	Huppertz et al.
5,758,257 A	5/1998	Herz
5,758,259 A	5/1998	Lawler
5,774,170 A	6/1998	Hite et al.
5,778,182 A	7/1998	Cathey et al.
5,794,210 A	8/1998	Goldhaber et al.
5,818,438 A	10/1998	Howe et al.
5,838,314 A	11/1998	Neel
5,848,396 A	12/1998	Grace
5,861,906 A	1/1999	Dunn
5,892,508 A	4/1999	Howe
5,892,536 A	4/1999	Logan
5,917,481 A	6/1999	Rzeszewski
5,948,061 A	9/1999	Merriman et al.
5,973,683 A	10/1999	Cragun
5,983,227 A	11/1999	Nazem
6,002,393 A	12/1999	Hite
6,005,597 A	12/1999	Barrett
6,029,045 A	2/2000	Picco et al.
6,076,094 A	6/2000	Cohen
6,100,916 A	8/2000	August
6,119,098 A	9/2000	Guyot et al.
6,172,674 B1	1/2001	Etheredge et al.
6,185,614 B1	2/2001	Cuomo
6,202,210 B1	3/2001	Ludtke
6,252,586 B1	6/2001	Freeman
6,282,713 B1	8/2001	Kitsukawa
6,314,568 B1	11/2001	Ochiai
6,314,577 B1	11/2001	Pocock
6,327,574 B1	12/2001	Kramer et al.
6,345,256 B1	2/2002	Milsted et al.
6,434,747 B1	8/2002	Khoo et al.
6,446,261 B1	9/2002	Rosser

6,457,010 B1 9/2002 Elderling et al.

6,463,468 B1 10/2002 Buch

6,463,585 B1 10/2002 Hendricks

6,477,704 B1 11/2002 Cremia

6,487,538 B1 11/2002 Gupta et al.

6,496,818 B1 12/2002 Ponte

6,505,348 B1 1/2003 Knowles

6,507,839 B1 1/2003 Ponte

6,510,417 B1 1/2003 Woods et al.

6,631,360 B1 10/2003 Cook

6,714,992 B1 3/2004 Kanodia et al.

6,718,551 B1 4/2004 Swix et al.

6,762,168 B2 7/2004 Chu

6,845,396 B1 1/2005 Kanodia et al.

6,850,988 B1 2/2005 Reed

6,861,618 B2 3/2005 Binet et al.

2002/0123928 A1 9/2002 Elderling et al.

2002/0143647 A1 10/2002 Headings

2002/0143847 A1 10/2002 Headings et al.

2002/0199197 A1 12/2002 Winter

2003/0067554 A1 4/2003 Klarfield

FOREIGN PATENT DOCUMENTS

EP	1162840 A2	12/2001
WO	97/17609	8/1994
WO	WO 99 04561 A	1/1999
WO	WO 99 45702 A	9/1999
WO	WO 99 52285 A	10/1999

OTHER PUBLICATIONS

Cauley, Leslie, "Microsoft Baby Bell Form Video Alliance," *The Wall Street Journal*, Sep. 28, 1994.

"allNetDevices—Report: Interactive TV Soon to Become Direct Marketing Tool," allNetDevices, www.devices.internet.com/com_cgi/print/print.cgi?url=http://devices..report_interactive.html, (Nov. 16, 2001) pp. 1-2.

"Liberate Technologies—Solutions," Digital Cable Two-Way, www.solutions.liberate.com/architecture/dc2.html, (Nov. 16, 2001) p. 1 of 1.

"Liberate Technologies—Solutions," Digital Broadband Telco, www.solutions.liberate.com/architecture/db.html, (Nov. 16, 2001) p. 1 of 1.

"Liberate Technologies—Solutions," Liberate Connect Suite, www.solutions.liberate.com/products/connect_suite.html, (Nov. 16, 2001) pp. 1-2.

"Liberate Technologies—Solutions," Liberate imprint Server™ www.solutions.liberate.com/products/imprint_server.html, (Nov. 16, 2001) p. 1 of 1.

"Liberate Technologies—Solutions," Liberate Mediicast Server™ www.solutions.liberate.com/products/mediicast_server.html, (Nov. 16, 2001) pp. 1-2.

"Spike High Performance server Array," Mixed Signals Technologies, Inc., www.mixedsignals.com.itvinfo@mixedsignals.com (2000) p. 1.

"power, flexibility, and control," RespondTV, www.respondtv.com/whyrespond.html (Nov. 16, 2001) pp. 1-2.

"It just clicks!" Respond TV, www.respondtv.com/inaction.html, (Nov. 16, 2001) pp. 1-2.

"The Wink System," Wink System Diagram, www.wink.com/contents/tech_diagram.html, (Nov. 16, 2001) p. 1 of 1.

"What is Wink?", www.wink.com/contents.whatiswink.html, (Nov. 16, 2001) p. 1 of 1.

"How Wink Works," What is Wink: How wink works, www.wink.com/contents/howitworks.html, (Nov. 16, 2001) p. 1 of 1.

"What is Wink: Examples," What is Wink—Examples, www.wink.com/contents/examples.html, (Nov. 16, 2001) pp. 1-2.

US 6,983,478 B1

Page 3

"Nielsen Media Research—Who We Are & What We Do," www.nielsenmedia.com/whoweare.html, (Oct. 11, 2001) pp. 1-4.

www.actv.com, Screen Print, Oct. 8, 2000.

"ACTV, Inc. Offers Exclusive Preview of 'Individualized Television' at Official All-Star Cafe," Business Wire, Sep. 16, 1998.

"ACTV's HyperTV & 'Individualized Television' to be Powered by Sun Microsystems JavaTV Technology," Business Wire, Apr. 21, 1999.

Whitaker, Jerry, "Interactive TV: Killer App or Technical Curiosity?", Broadcast Engineering, Dec. 1999.

Dickson, Glen, "Digital TV gets specifically directed," Broadcasting & Cable, Jun. 5, 2000.

Reed, David, "The future is digital," Precision Marketing, v. 13, n.51, p.27, Sep. 21, 2001.

U.S. Appl. No. 09/799,306, filed Jan. 8, 1997, Grauch.

U.S. Appl. No. 10/016,988, filed Dec. 14, 2001, Matz.

U.S. Appl. No. 10/017,111, filed Dec. 14, 2001, Matz.

U.S. Appl. No. 10/017,630, filed Dec. 14, 2001, Swix.

U.S. Appl. No. 10/017,640, filed Dec. 14, 2001, Matz.

U.S. Appl. No. 10/017,742, filed Dec. 14, 2001, Matz.

U.S. Appl. No. 10/036,677, filed Dec. 21, 2001, Swix.

U.S. Appl. No. 10/036,923, filed Dec. 21, 2001, Swix.

U.S. Appl. No. 10/037,005, filed Dec. 21, 2001, Matz.

U.S. Appl. No. 10/778,345, filed Feb. 17, 2004, Swix.

* cited by examiner




Figure 1




Figure 2




Figure 3




Figure 4A




Figure 4B




Figure 5




Figure 6A

1 damage and irreparable injury unless and until TiVo is enjoined by this Court from continuing
2 such infringement.

3 **COUNT THREE: INFRINGEMENT OF U.S. PATENT NO. 6,118,976**

4 21. AT&T incorporates by reference Paragraphs 1 through 20, as if fully set forth
5 herein.

6 22. United States Patent No. 6,118,976, entitled "Asymmetric Data Communications
7 System" (hereinafter, the "'976 patent"), duly and legally issued on September 12, 2000 after a
8 full and fair examination. AT&T Intellectual Property I, L.P. is the assignee of all rights, title,
9 and interest in the '976 patent, including the right to sue and recover for all past infringement. A
10 true copy of the '976 patent is attached as Exhibit C.

11 23. As an example, and not to be limited to only a single infringing product, TiVo has
12 developed, has tested, promotes, markets, and sells DVR units that allow for user interaction
13 through a communications channel ("interactive DVR units"). Such products (or the use or
14 operation of such products) fall within the scope of one or more claims of the '976 patent.

15 24. TiVo has infringed and continues to infringe the '976 patent, by, among other
16 acts, making, using, offering for sale, selling, and/or importing within this Judicial District and
17 elsewhere in the United States, without license or authority by AT&T, products and/or processes
18 covered by one or more claims of the '976 patent, including, but not limited to, interactive DVR
19 units.

20 25. As a consequence of TiVo's infringement, AT&T is entitled to recover damages
21 adequate to compensate it for the infringement complained of herein, but in no event less than a
22 reasonable royalty.

23 26. TiVo has caused and will continue to cause AT&T substantial damage and
24 irreparable injury by virtue of its past and continuing infringement of the '976 patent. TiVo will
25 continue to infringe the '976 patent unless enjoined by this Court. AT&T will suffer further
26 damage and irreparable injury unless and until TiVo is enjoined by this Court from continuing
27 such infringement.




Figure 6B

Clickstream Data 80

	Timestamp	Application ID	# Bytes to Follow	Application Specific Data	
				Event ID	Channel ID
Event Record	3:59:30pm	Cable App.	4 bytes	Power On	ABC
Event Record	4:00:00pm	Cable App.	4 bytes	Channel Up	NBC
Event Record	4:04:17pm	Cable App.	4 bytes	Channel Up	TBS
Event Record	4:06:25pm	Cable App.	4 bytes	Channel Dwn	NBC
Event Record	4:15:45pm	Cable App.	4 bytes	Channel Up	TBS
Event Record	4:55:45pm	Cable App.	4 bytes	Power Off	None

Prevue Guide Data 82

	Content ID	Channel ID	Start Time	End Time
Content Record	National News	ABC	3:30:00pm	4:00:00pm
Content Record	Murphy Brown	NBC	3:59:00pm	4:30:00pm
Content Record	Simpsons	NBC	4:30:00pm	4:59:00pm
Content Record	N.G.Explorer	TBS	3:05:00pm	4:05:00pm
Content Record	Andy Griffith	TBS	4:05:00pm	4:35:00pm
Content Record	NBA Basketball	TBS	4:35:00pm	6:05:00pm

Broadcast Advertising Data 84

	Content ID	Channel ID	Start Time	End Time
Content Record	Coca Cola #10	NBC	4:00:30pm	4:01:00pm
Content Record	Visa #2	NBC	4:01:00pm	4:01:30pm
Content Record	Delta #1	TBS	4:03:30pm	4:04:00pm
Content Record	Delta #3	TBS	4:04:00pm	4:04:30pm
Content Record	Visa #21	TBS	4:04:30pm	4:05:00pm

Clickstream Timeline 92

	Timestamp	App. ID	# Bytes to Follow	Application Specific Data		
				Event ID	Channel ID	Content ID
Event Record	3:59:30pm	Cable App.	6 bytes	Power On	ABC	National News
Event Record	4:00:00pm	Cable App.	6 bytes	Channel Up	NBC	Murphy Brown
Event Record	4:00:30pm	Cable App.	6 bytes	Change Content	NBC	Coca Cola #10
Event Record	4:01:00pm	Cable App.	6 bytes	Change Content	NBC	Visa #2
Event Record	4:01:30pm	Cable App.	6 bytes	Change Content	NBC	Murphy Brown
Event Record	4:03:17pm	Cable App.	6 bytes	Channel Up	TBS	N.G.Explorer
Event Record	4:03:30pm	Cable App.	6 bytes	Change Content	TBS	Delta #1
Event Record	4:04:00pm	Cable App.	6 bytes	Change Content	TBS	Delta #3
Event Record	4:04:30pm	Cable App.	6 bytes	Change Content	TBS	Visa #21
Event Record	4:03:17pm	Cable App.	6 bytes	Change Content	TBS	N.G.Explorer
Event Record	4:05:00pm	Cable App.	6 bytes	Change Content	TBS	Andy Griffith
Event Record	4:06:25pm	Cable App.	6 bytes	Channel Dwn	NBC	Murphy Brown
Event Record	4:15:45pm	Cable App.	6 bytes	Channel Up	TBS	Andy Griffith
Event Record	4:35:00pm	Cable App.	6 bytes	Change Content	TBS	NBA Basketball
Event Record	4:55:45pm	Cable App.	6 bytes	Power Off	None	

Figure 7

METHOD AND SYSTEM FOR TRACKING NETWORK USE

This application is a continuation of U.S. application Ser. No. 08/779,306, filed with the U.S. Patent Office on Jan. 6, 1997, now abandoned.

This invention is a method and system for tracking subscriber use of network applications, particularly network applications involving delivery of interactive media or video programming.

BACKGROUND OF THE INVENTION

Broadcast and cable television have long dominated the visual media market. New communications technologies, however, have accelerated demand for new types of media such as video on demand, interactive video, interactive gaming, home shopping or interactive advertising. Unlike broadcast television, viewers of these services typically are paying "subscribers," although payments from advertisers also pay a large share of the costs of providing these media services.

To gauge the effectiveness of their spending, advertisers have long sought information on viewers' viewing patterns. A number of devices and techniques exist for gathering such information. For instance, U.S. Pat. No. 4,258,386 to Cheung and U.S. Pat. No. 4,556,030 to Nickerson, et al., describe the general concept of deploying in viewers' homes devices for monitoring a viewer's television set ("TV") in order to accumulate data illustrating viewing habits such as which channels were watched at particular times. Accumulated data is then forwarded via telephone lines to a central location for analysis. Cheung sends data from particular monitoring stations at a preselected, specific "window" of time; interruptions to transmission during that window result in the Cheung system forwarding the data at another time.

Other systems and methods provide somewhat more use data than just channel numbers viewed and time of viewing. Typically, however, the information is for a smaller subset of users. Thus, U.S. Pat. No. 4,816,904 to McKenna, et al., U.S. Pat. No. 4,912,552 to Allison, III, et al. and U.S. Pat. No. 5,374,951 to Welsh, all disclose monitoring "panelist" TV use in order to collect data about panelist viewing patterns as well as certain marketing information. Generally, panelist monitoring is used to gauge the effectiveness of advertising on selected groups of "panelists," each of which is one household in a group comprising a "panel," typically located in a particular geographical area.

Monitoring not only determines which commercial and TV programs the panelist views but also may be used to gather information about which products panelists purchase. For instance, the U.S. patent to McKenna discloses a remote data collection unit located at a panelist home that monitors viewer identification data and TV functions (e.g., channel viewed, VCR viewing time or game time). Additionally, a wand is provided for inputting bar codes of purchased items. Monitored data is sent via the telephone network to a central location, which can also download questionnaires to the panelist and receive responses. Allison and Welsh disclose similar monitoring systems and methods. Instead of simply monitoring the channel number that a panelist was viewing at a particular time, Welsh discloses monitoring identification information carried in the television signal vertical blanking interval that identifies preselected commercials. After detecting and storing the identification information

that identifies particular commercials viewed by panelists, the data is transmitted by telephone to a central location for analysis.

Monitoring systems also have been used with some early interactive media systems. U.S. Pat. No. 5,404,393 to Remillard discloses an interactive TV system. Among other elements of the system, a controller monitors TV channels and time/date stamps the selected channel so that, indirectly, viewers' programming choices may be monitored. Data is assembled into a "user profile," which is uploaded to an appropriate facility via the telephone network.

Nevertheless, while panelist monitoring systems like those of Allison, McKenna and Welsh or interactive television monitoring systems like Remillard's provide somewhat more monitoring data than just TV tuning data, they do so only for limited groups. For example, when more data is gathered (like purchase information), it is done only for the panelist groups, rather than for subscribers to the entire system. Also, systems like McKenna's that uses a wand for scanning bar codes are intrusive systems that require user action to collect data rather than collecting data passively and automatically. Other systems contemplate capturing only some of the data generated by subscriber's viewing activities or only some of the ratings information. For instance, previous systems typically capture ratings information that identify television shows viewed rather than whether the subscriber viewed commercials displayed during those shows.

Perhaps more importantly, none of the systems described attempt to match "raw" information on channels viewed with programming information. Nor do those systems match viewing pattern information with demographics information about the particular users in order to provide more "targeted" advertising.

SUMMARY OF THE INVENTION

The present invention uses a collector, associated with a subscriber's set top box ("STB"), to obtain data about any "events"—subscriber actions or changes in programming—that are of interest. Data about virtually any events, from channels watched to volume changes to interactive applications invoked, may be captured with the collector. Event records comprising such data, as well as the identity of the application involved and the event time, are buffered. Periodically or on command, event records are uploaded from the buffer to a merge processor such as through an interactive network that allows for duplex communication with the STB. The merge processor, which may be a head end server or a workstation computer forming part of or coupled to the media delivery network, receives (1) the event data and (2) content data that identifies programming content broadcast or delivered throughout the region in which the system is deployed. Timelines showing particular events over time may then be generated for each subscriber. Rather than just determining the channel viewed and time of day, the event timelines describe the programming or interactive applications selected by or shown to a subscriber over a selected period of time (e.g., 24 hours).

The merge processor may further filter this collected and merged data to generate reports ranging from descriptions of a single user's viewing patterns to very high level viewing patterns showing the number of users who watched or participated in a particular program for a selected time period. Further, that information can be combined with billing and demographics information to provide detailed

information on a particular subscriber's or group of subscribers' viewing and related buying patterns.

The present invention thus involves a method for obtaining detailed information on every application invoked by a subscriber and information about the type of programming shown. The first step is to identify data that describe the events of interest that occur. Those events include: the channel viewed, a switch to another channel, a passive change in programming because of a commercial break or change to a new program, use of a VCR or other ancillary device, or invocation of an interactive application and subscriber commands given to the system during the application. Event data also includes start and stop times, identification of the subscriber's STB or specific data needed to be recorded for any particular interactive or other application.

Event records are formed from this collected data and buffered before uploading through the interactive or other media delivery network to a headend, server or third party data analysis system. Before uploading, the captured data may be compressed and formed into packets for transmission.

Using the system or method of the present invention allows service providers to obtain ratings information and detailed information on subscriber viewing patterns and subscriber use of interactive applications. Thus, the present invention can track the number of subscribers viewing or watching particular programs, including advertisements. It also can track use of particular interactive applications such as video on demand. The invention automatically matches data describing programming content with event data describing a channel or application activated or controlled by the subscriber. This allows the invention comprehensively to track user "channel surfing." Also, the invention can compare subscriber demographics or billing information with viewing pattern information in order to tailor commercials to those subscribers; determine whether subscribers with a selected demographic background viewed a commercial of interest; or determine the demographics of subscribers that viewed selected commercials.

Persons skilled in the art will recognize that the present invention may be used with numerous types of networked media delivery systems. For instance, the method or system of the present invention can be deployed on an interactive media delivery system or modified for use with a conventional cable television network, a wireless cable television network, or a home satellite television network.

It is accordingly an object of the present invention to provide a system and method for collecting information about patterns of subscriber viewing and use of a media delivery system.

It is another object of the present invention to provide a system and method for determining which network applications, particularly interactive applications, are invoked by particular subscribers.

It is an additional object of the invention to provide a system and method for communicating collected information to a merge processor.

It is a further object of the invention to provide to the merge processor information about the programming content distributed over the media delivery system.

It is yet another object of the invention to provide a system and method for merging the collected information with the programming information in order to obtain comprehensive information about programming shown to or network applications invoked by subscribers. Other objects, features and advantages of the present invention will become apparent upon reading the rest of this document.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows a block diagram of elements of one embodiment of the system of the present invention.

FIG. 2 shows a block diagram of a Set Top Box as used with the embodiment of the present invention shown in FIG. 1 and provided with a clickstream processor.

FIG. 3 shows a schematic diagram showing the upload cycle for collected event data.

FIGS. 4A and 4B show the upload of collected event data from a selected Set Top Box through the network to the staging server shown in FIGS. 1 and 5.

FIG. 5 shows an overview of the staging server, its functions and its interconnections with various data sources.

FIG. 6A shows the system elements required for merging and parsing the event and content data collected by the present invention.

FIG. 6B shows the assignment of priority to content data necessary for completing the merge and parse process.

FIG. 7 shows the results of a merge and parse process of the present invention.

DETAILED DESCRIPTION OF THE DRAWINGS

System Overview

FIG. 1 shows a block diagram of the components of the system 20. System 20 is a demographics and programming ratings collection and analysis system that may be deployed for use on an interactive media delivery system such as the Interactive Video Services Network deployed by BellSouth Interactive Media Services. That interactive system is described in co-pending application Ser. No. 08/428,718, assigned to the assignee of the present invention and which document is hereby incorporated in its entirety by this reference. However, persons skilled in the art will recognize that the present invention may be used with any of a variety of interactive media delivery systems, standard or wireless cable television systems, satellite television systems or other media delivery systems that allow duplex communication (perhaps with the return path via a separate (e.g., telephone) network) to a set top box ("STB") 30 coupled to a subscriber's display device, such as a television set or the like.

In any event, FIG. 1 shows various system 20 elements and subsystems that communicate with each other to transmit collected information, data error detection schemes and data acknowledgments. Briefly, the STB 30 communicates through a distribution network 52 with a video server 60, such as a video transfer engine ("VTE"), that may be acquired from Hewlett Packard ("HP"), with a video/object storage database 54. Video server 60 couples to a video control server 56, such as an Inter Media Server available from Sybase and deployed on a platform such as an HP 9000, with a database 58. The video server control 56 controls video server 60 and also logs information about video server 60 use. A staging server 70 receives collected records of events of interest. These "event records" pass through the video server control 56, which also couples to a Marketing and Information System ("MKIS") 100 that couples to staging server 60, which receives (1) the event records and (2) content data from various sources 120, 140 and 160 identified in FIG. 1 and which describe programming content available through the interactive network to all subscribers. MKIS 100 may be coupled to a third party search and analysis system 110 that can provide customer support operations.

STB 30 provides a platform by which (1) content is converted to a selected video format (e.g., NTSC or PAL) and presented to the subscriber or (2), for interactive systems, messages are exchanged (including video data) over a network 52 with the staging server 70. STB 30 also could include platforms capable of: (1) receiving messages from a user input device, such as a hand-held remote control unit; (2) translating video signals from a network-native format into a format that can be used by the television or display device; (3) inserting alphanumeric or graphical information into the video stream in order to "overlay" that information on the video image; (4) providing graphic or audio feedback to a user; or (5) possibly the most basic function, simply routing a traditional broadcast signal to a viewing device connected to the STB 30. Analogous terms to STB include: Set-Top Terminal ("STT"), Cable Converter, and Home Communications Terminal ("HCT") and any of these devices may be coupled to or made a part of a display device for showing programming to subscribers. Generally, STB 30 may be a Richmond or 8600X available from Scientific Atlanta, a CFT 2200 available from General Instruments, Thomson's DSS or any other device equipped with (1) a microprocessor; (2) memory for operating instructions and storage; and (3) a control interface for accepting subscriber commands from a remote control device or control panel.

For the particular embodiment of system 20 shown in the Figures, collected event records that are packaged for transport through system 20 are called "clickstream" data or information. FIG. 2 shows a clickstream processor 34 that resides in the memory, such as DRAM or the like, of an STB 30 and which has a clickstream kernel 36, buffers 42 or 44, a clickstream upload handler 40, a clickstream controller 38 and a clickstream event API (application programming interface) 41.

Briefly, the clickstream kernel 36 buffers events passed to it by various network applications through the clickstream event API 41. Clickstream controller 38 accepts control messages from staging server 70 and appropriately stores their payload. Typical messages may be sent over the Extended Super Frame ("ESF") pass-through data link and control the uploading of clickstream data. Clickstream upload handler 40 accepts control messages over the system 20, which messages control the uploading of collected clickstream data over the reverse path through network 52. Also, the clickstream upload handler 40 stores the payload of these messages in appropriate and available memory and accepts the messages sent to it to acknowledge the receipt of uploaded clickstream data.

Referring again to FIG. 1, video server 60 provides information from video/object storage 54 to the particular interactive system over which system 20 is deployed. Clickstream data collected at STBs 30 can be uploaded to staging server 70 in any number of ways. For instance, FIG. 1 shows that the distribution network 52 could couple directly to staging server 70, allowing clickstream data packets sent from STBs 30 to be forwarded to staging server 70 directly and for staging server 70 to then return via the network 52 data acknowledgements. A network management controller 50 controls the flow of information through the network 52. Alternatively, FIG. 1 and, in greater detail, FIG. 4B, show that clickstream data packets may be sent to the distribution network 52 to the video server 60. Video server 60 passes through both clickstream data uploads from various STBs 30 and data acknowledgments returned to the STBs 30. A communications router inside the video server 60 redirects traffic to the appropriate destination. Video server control 56 similarly acts as a pass-through device for STB 30 click-

stream data going to the staging server 70 and as a pass-through device for staging server 70 data acknowledgments to the STBs 30. Also, video server control 56 may provide log information that identifies interactive applications invoked by particular STBs 30. That log information is provided to the staging server 70 so that video server control 56 also acts as another data source about content available over the network, like EPG metadata source 120, broadcast advertising metadata source 140, or advertising traffic control metadata source 160. Staging server 70 collects all such clickstream data and content data, analyzes and then forwards it to MKIS database 100 or to a third-party analysis engine and database 110, as described in more detail in the text associated with FIGS. 5-7.

Journaling of Event Data

Clickstream processor 34 collects information to create a "journal" or log about all events or selected events of interest. An event is an action or a change in the state of a STB 30 that is deemed important to building a knowledge base on subscribers or their viewing patterns. For example, an event can include key presses to change channels or volume, mute, to enter the navigator for the interactive system, to turn the STB 30 off or on, to fast forward, to pause or to rewind a video obtained via the video on demand application. The events include applications called by the subscriber, such as interactive gaming applications, an electronic program guide, a video on demand or near video on demand application, a home-shopping application or a particular company's interactive application, such as The Weather Channel's weather on demand, World Span's travel on demand or Light Span's educational interactive application. Events include subscriber use of and control commands to peripheral devices coupled to the STB 30 or a subscriber's display device, such as a VCR or videodisk player.

Each application residing on the STB 30 interfaces with the clickstream processor 34 to send selected data for maintaining a desired journal. Assuming that the system 20 is used with an interactive system, many different applications may be deployed over that system and may be triggered by the subscriber. Some fairly typical applications that might be invoked include:

a cable television application that handles subscriber remote controls (like channel or volume changes);
an electronic programming guide application such as TV Data, Prevue or Star Sight interactive services;
an interactive game;
a video on demand or near video on demand application;
company specific applications, that might be offered by content provider such as the Weather Channel, MTV, Showtime, etc.; or
a navigator application to help the user choose options.

Each of these applications, as well as some internal applications that the system 20 may wish to monitor, will be assigned a unique application identifier.

Clickstream processor 34 interfaces with the various applications resident in the STB 30's operating system 32 and any third party applications 33. Note that for systems using other types of STB 30's than the embodiment described in the Figures, those STB 30's need not have an operating system. Instead, all instructions can be written directly to the memories of those particular STBs. Applications 33 can be added by either downloading entirely new software directly to memory or by downloading new tables as described below.

When an application 33 reaches a point where an "event" of interest has been generated, the application 33 stores an event record to memory. The application 33 then launches to the clickstream kernel 36 the event record, including information such as: (1) the application's 33 identification code (e.g., the "Cable Television Application" or a particular interactive application); (2) a count of the amount of information (number of bytes) to be journaled; (3) a "time stamp" that defines a unique point in time, e.g., by defining the date and time of day, accurate to the hour, minute or second; (4) an identification code for the event, or (5) where the event data was stored. Clickstream kernel 36 uses the information provided by the applications 33 to collect the event data, format it and place it into a buffer 42 or 44. Table I shows the type of information that will be generally sent by the clickstream processor 34 to the buffers 42 or 44.

TABLE I

<u>Application Event Record</u>		20	Size
Timestamp	6 bytes		
Assigned Application ID	16 bits		
Number Bytes to Follow (length)	8 bits		
Application Specific Data with customized formats and lengths	Multiple Bytes	25	

Global table II defines events of interest that each application can identify, collect, store in the "Application Specific Data" field and notify the clickstream kernel 36. These events could be as simple as a broadcast channel change by pressing the "Chan Up" remote key. All of these event types can be accessed and used by each application. While each application may not use every possible event type, the number of events available for collection allows system 20 to extract any pertinent usage information for analysis. Also, the use of the global table II increases system 20 efficiency because event types can be modified, added or removed.

TABLE II

<u>EVENT DEFINITIONS</u>		40	45
Code	Event		
<u>Content Related Events</u>			
0x0000	Passive Content Change		
	Direct Key Presses		
0x0001	TV < > ITV		
	Pressed		
0x0002	Power Pressed		
0x0003	One (1) Pressed		
0x0004	Two (2) Pressed		
0x0005	Three (3) Pressed		
0x0006	Four (4) Pressed		
0x0007	Five (5) Pressed		
0x0008	Six (6) Pressed		
0x0009	Seven (7) Pressed		
0x000A	Eight (8) Pressed		
0x000B	Nine (9) Pressed		
0x000C	Zero (0) Pressed		
0x000D	Channel Up Pressed		
0x000E	Channel Down Pressed		
0x000F	Volume Up		

Note that Table II defines relative volume changes (e.g. "volume change below 50%," "volume change below 25%," etc.). Although the applications could capture the actual key presses that lead to these relative volume changes, that level of detailed information is of little use to system 20 operators. Also, capturing all that detail leads to more records and higher demands upon the transmission network 52 when those records are uploaded. Applications could also be configured to "filter" other unwanted details about other subscriber activities. For example, when subscribers "channel surf" by quickly flipping through a number of channels in a short period of time, the application could be configured not to record channel changes unless the subscriber paused for greater than a certain selected time period (e.g., 15 to 30 seconds). Again, this eliminates information of little use and decreases network traffic.

Table III defines a small portion of a sample global channel identification table that proposes codes for identifying national and local broadcasters. Such a table allows any application journaling events which occur while subscribers are viewing broadcast or cable television programs

to identify the network carrying the programming content by using a subset of the global table II. In this way channel lineups can be changed yet the identifier for a broadcast or cable network would stay the same. The use of this mapping scheme eliminates the need to map an ever-changing channel number to a network.

TABLE III

Broadcast Channel Identification		10
0x0100 to 0x011F	News/Talk Shows	15
0x0100	CNN	15
0x0101	Headline News	
0x0102	The Weather Channel	
0x0103	CNBC	20
0x0104	CSPAN	
0x0105	CSPAN-2	
0x0106	America's Talking	
0x0107	Talk Channel	
0x0108	Court TV	
0x0109	The Crime Channel	25
0x010A	National Empowerment TV	
0x0120 to 0x013F	Sports	30
0x0120	ESPN	
0x0121	ESPN-2	
0x0122	SportSouth	
0x0123	The Golf Channel	
0x0124	Classic Sports Network	
0x0125	Prime Network	35
0x0126	NewSport	
0x0140 to 0x015F	Music	40
0x0140	MTV	
0x0141	VH-1	
0x0142	Country Music Television	
0x0143	The Nashville Network	45
0x0144	The Box	
0x0145	Video Jukebox	
0x0146	MOR Music TV	
0x0147	Music Choice	50

Table IV below shows some possible identification codes for particular applications. Note that each application could be programmed to insert its application ID code into the event record without accessing table IV. But by having each application access the table IV during the journaling process, the system's 20 ability to modify or add application ID codes easily is enhanced because such codes could be populated across system 20 by downloading an updated table IV. Providing for downloading of new tables increases the application footprint and system 20 complexity so tables can also be part of the application programming.

TABLE IV

Application Identifiers	
ID Code	Content
0x0000	Operating System
0x0001-F	Operating System Sub-Systems
0x0010	Application Manager
0x0011	Cable Television Application
0x0012	Clickstream Kernel
0x0100	EPG System
0x0101	Digital Pictures - Interactive Game
0x0110-F	Viacom - MTV>Showtime, etc.
0x1000	Interplay Written Applications General ID
0x1001	Interplay Runtime Engine
0x1002	Interplay Navigator
0x1003	Interplay VOD
0x1004	Interplay NVOD
0x1005	Interplay TownGuide
0x1100	The Weather Channel, Weather On-Demand
0x1101	Worldspan - Travel On-Demand
0x1102	Lightspan - Educational Interactive Application
0xFFFF	Missed Events Record

25 Each particular application can simply reference the global application, event and channel identification tables (which periodically may be updated and then downloaded to STBs 30) in order to build an event record. Examples of application specific event records that may be created in this manner are shown in Tables V through VIII below and discussed in associated text.

30 35 A cable TV application 33 may tune analog or digital broadcast services. When a command to change channels is entered, the cable TV application 33 is invoked. The cable TV application 33 begins building an event record by inserting an application ID and time stamp into the record.

40 Next, the application 33 determines the "event ID" by cross-referencing the command with the global event ID table II for the proper code. Then, the application 33 journals the "Channel ID."

45 45 Although the Channel ID could simply be the number of the channel, that information means little. The fact that channel 6 was watched more than channel 7 has little or no meaning unless networks and, ultimately, the content delivered by those networks are associated with particular channels. Accordingly, the Channel ID may be a field, like a 16 bit field, which uniquely identifies the broadcast network displayed on that particular channel. The Channel ID may be determined by programming the cable TV application 33 to

50 55 compare the channel number tuned with global broadcast channel identification table II, above, to determine the correct channel identification code. Correlating the channel number with the channel identification code found in Table III ensures accurate reporting even though channels may differ at different cable TV headends within a particular region or even though individual channel line-up changes may be made over a period of time. This correlation between channel number and channel identification code could be done also at the staging server 70 after it receives all of the event records, provided that correlation there accounted for different regional channel lineups.

TABLE V

Cable TV Application Event Record

	Size	5
Application ID: See Application ID table IV	16 bits	
Timestamp: Identifies event time	6 bytes	
Event ID: See Global Event ID table II for Syntax	16 bits	
Channel ID: See Broadcast Channel ID table III for Syntax	16 bits	

Table VI below shows a navigator application that may be provided in order to give subscribers an interactive menu that assists them in selecting from the many available programs and applications in an interactive network. The "Event ID" refers to the identification codes for commands relating to the Navigator application, which codes may be located by referring to the global event ID table II above. Table VI also shows some of the features of the navigator that might be used by the subscriber and that could be useful to track. The right hand column under "Size/Data" shows, first, next to the "Application state ID" that 8 bits are allocated to that record and, second, in the various rows beneath, the particular code that is journaled in order to indicate a subscriber accessed the identified (e.g. Fly-Thru, Main Menu, etc.) screen. Such information lets system 20 operators determine the screens that users are viewing heavily or lightly in order to replace less popular screens with more useful ones or to charge more for advertisements placed on heavy use screens.

TABLE VI

Navigator Application Event Record

	Size/Data	
Application ID: See Application ID table IV	16 bits	
Timestamp: Identifies event time	6 bytes	
Event ID: See Global Event ID table for Syntax	16 bits	
Application State ID: See below for information tracked:	8 bits	
Fly-Thru	0x00	
Main Menu	0x01	
Information (Help) Screen or Video	0x02	
Movies Sub-Menu	0x03	
Movie Categories Sub-Menu	0x04	
List of Movies Sub-Menu	0x05	
Movie Info Screen	0x06	
Movie Buy State	0x07	

Table VII similarly shows the journaling information collected for a video on demand application 33 that may be launched in an interactive service from the Navigator application above or its equivalent. Some of the information collected here may include the amount of pausing, fast forwarding and rewinding. Additionally, the service provider may want to determine whether viewers are recording a video in order to charge them a recording fee. Similar information could be collected for a near video on demand service, which typically allow only incremental pause, forward or rewind.

TABLE VII

Video on Demand Application Event Record

	Size/Data	
Application ID: See Application ID table IV	16 bits	
Timestamp: Identifies event time	6 bytes	

TABLE VII-continued

Video on Demand Application Event Record

	Size/Data	
Event ID: See Global Event ID table for Syntax	16 bits	
Application State ID: See below for information tracked:	8 bits	
Playing	0x00	
Paused	0x01	
Fast Forward	0x02	
Rewind	0x03	
Info (Help) Video or Screen Played	0x04	
reserved	0x05	
reserved	0x06	
reserved	0x07	

Table VIII below shows the event record for the Electronic Program Guide (EPG) application 33. The EPG application 33 records the application ID, timestamp and event ID records just as do the above applications described in tables V-VII. Additionally, it has an application 33 state ID field that identifies which of the display screens were accessed by subscribers, as shown below.

TABLE VIII

Electronic Program Guide (EPG) Application Event Record

	Size/Data	
Application ID: See Application ID table IV	16 bits	
Timestamp: Identifies event time	6 bytes	
Event ID: See Global Event ID table for Syntax	16 bits	
Application State ID: See below for information tracked:	8 bits	
Initial Display Screen	0x00	
Look Ahead Display 4 Hour	0x01	
Look Ahead Display 8 Hour	0x02	
Look Ahead Display 12 Hour	0x03	
Look Ahead Display 16 Hour	0x04	
Look Ahead Display 20 Hour	0x05	
Look Ahead Display 24 Hour	0x06	
Reserved	0x07	

Generally, similar information about other applications 33, such as home shopping, interactive gaming or any other new applications deployed over an interactive or other media delivery system, can be tracked in a similar fashion. Additionally, the journaling process may be used to track errors within the system 20, with clickstream kernel 36 journaling such errors using the same method as described above.

Over time, the journaling needs of system 20, or system 20 itself may evolve. Applications may be changed or new ones deployed. New events may become of interest to the operator of system 20. In order to provide flexibility for system 20, operators may download to STBs 30 new or replacement applications that will include the necessary processes for journaling all events of interest.

Sample Journal

Assume that Mr. Smith turns on his interactive television at 7:30 p.m. to watch a half hour news program on channel 5, which corresponds to CNN for that region. At 8:00 p.m. he accesses the Navigator application to order a video through the video on demand application. He then accesses the Video on Demand application, which automatically begins playing a video at 8:04, pauses the video at 8:50 and

begins playing again at 8:55 until it is completed at 9:45, at which point he turns off his interactive TV.

Mr. Smith's activities generate the following event records shown in table IX below (for convenience, multiple events occurring under a single application are grouped even though separate records are created in operation):

TABLE IX

<u>Sample Event Records</u>		Data
<u>Cable Application Event Record 1 Content</u>		
Application ID: See table IV for application ID Code	0x0011	
Timestamp: Identifies event time	1/1/96	
	7:30:01 p.m.	
Event ID: See Global Event ID table II to retrieve code for "power pressed"	0x002	
<u>Cable Application Event Record 2 and 3 Content</u>		
Application ID: See table IV for application ID Code	0x0011	
Timestamp: Identifies event time (Date will be same for second entry)	1/1/96	
	7:30:03 p.m.;	
	8:00:01 p.m.	
Event ID: See (1) global Broadcast Channel ID table III to determine that Channel 5 is CNN and locate code and (2) table II for an event ID code corresponding to an "ITV Press" by Mr. Smith.	0x0100	
Navigator Application Event Record 4 Content	0x0001	
<u>Application ID: See table IV for application ID Code</u>		
<u>Timestamp: Identifies event time for accessing each screen.</u>	0x1002	
Event IDs: See table II for event ID code that identifies an "enter" command by Mr. Smith to invoke this application.	1/1/96	
	8:01:30 p.m.	
Application State ID Code: This shows Mr. Smith accessed the Main Menu	0x0021	
Navigator Application Event Records 5-6 Content	0x01	
<u>Application ID: See table IV for application ID Code</u>		
<u>Timestamp: Identifies event time for accessing each screen. A separate record is created for each activity, with a timestamp showing initiation of each activity. Each record will have the corresponding event and state.</u>	0x1002	
Event IDs: See table II for event ID code that identifies an "enter" command by Mr. Smith to invoke this application.	1/1/96	
	8:02:00 p.m.;	
Application State ID Codes: This shows Mr. Smith accessed the Movies Sub-Menu and Movie Sub-menu list.	8:03:00 p.m.;	
Video on Demand Application Event Records 7-9 Content	0x0021	
Application ID: See Application ID table IV (same for each record).	0x0021	
Timestamp: Identifies event time for each event recorded by the application. (The day is the same for each record and each time corresponds with the activity identified below).	0x03	
Event ID: See table II for event ID codes that identify Mr. Smith's play, pause and play commands.	0x05	
Application State ID Codes: These show Mr. Smith activated this application, played, paused and then played again his selected video.	0x1003	
Cable Application Event Record 10 Content	1/1/96	
Application ID: See table IV for application ID Code	8:04:00 p.m.	
Timestamp: Identifies event time	8:50:00 p.m.	
	8:55:00 p.m.	
Event ID: See Global Event ID table II to retrieve code for "power pressed"	0x0022	
	0x0024	
	0x0022	
Application State ID Codes: These show Mr. Smith activated this application, played, paused and then played again his selected video.	0x00	
Cable Application Event Record 10 Content	0x01	
Application ID: See table IV for application ID Code	0x00	
Timestamp: Identifies event time	0x00	
Event ID: See Global Event ID table II to retrieve code for "power pressed"	0x002	

Event Record Upload Cycle

The variably sized event records are collected and then stored in one of two clickstream buffers 42 or 44. Capacity of each of the buffers may be statically provisioned or the system 20 may addressably download to particular STBs 30 an appropriate buffer 42 or 44 size. A buffer 42 or 44 may be an allocated, contiguous free area of STBs' 30 memory set aside for buffering event records only. Although advanced database techniques like link lists or record pointers could be used, they would increase the application footprint and complexity. Because buffer sizes of about 15 kB would probably accommodate the journaling needs of most applications, advanced database techniques need only be used for larger buffers. Buffers up to 15 kB should allow at least 4 to 8 hours of peak channel "surfing" between uploads (channel surfing typically will generate the most event records). In any event, empirical analysis of network use should determine the optimum buffer size.

Event records are directed to one of the two buffers 42 or 44, although a single or even more buffers could be used with the system 20. Conceivably, the system 20 could also be modified to upload event records in real time; however, that severely increases the possibility of instantaneous overloads in network traffic. Thus, system 20 preferably uses buffers 42 or 44 to buffer collected event records until they are upload.

Event records from a particular STB 30 may be uploaded in a format that assists in their transmission back through the distribution network 52 to the staging server 70. A header record may indicate the time the buffer 42 or 44 was first opened, the number of bytes in the buffer 42 or 44, the originating STB 30 by address, the version of the clickstream kernel 36 which generated the record and the type of data compression used on the following data (if any). This first header record may be of fixed length and uncompressed. Information following "Compression Type" may be compressed to save in transmission bandwidth. Table X below shows this general header format:

TABLE X

<u>Buffer Header Record</u>		Size
45	Transaction Code	8 bits
	Clickstream Version Number	8 bits
	Timestamp	6 bytes
	Number of Bytes	8 bits
50	STB Unique Address Most Significant	16 bits
	STB Unique Address Least Significant	32 bits
	Compression Type	8 bits

When (1) a buffer 42 or 44 fills, (2) an upload timer event expires or (3) upon command from the staging server 70, the clickstream processor 34 initiates an upload process. During that process the uploading buffer 42 is locked and subsequent event records are routed to and stored in the second buffer 44. When upload of buffer 42 is completed, records continue to buffer 44 until the next upload time, after which buffer 44 locks and records go to buffer 42. This cycle continues to repeat.

FIG. 3 shows an upload cycle diagram illustrating one method of evenly distributing increased traffic on the network 52 caused by upload of event records. The clickstream upload cycle consists of several parameters that define a start time and a cycle over which the uploading of data occurs. The "first occurrence" parameter defines a starting time in

history from which the cycle runs. The "cycle time" parameter defines the amount of time which elapses between periods of the upload cycle. When a cycle is complete the "upload duration" time starts, and the clickstream processor 34 of each STB 30 will randomize an exact upload time within the upload duration. This timing of upgrades will distribute the network load evenly over the entire upload duration period. 5

An example of the use of these parameters would be to define a period of time every day for STBs 30 within system 20 to upload data. Typically, the system 20 operator may want the data available every morning for analysis. Peak use of broadcast prime time or of interactive services will typically be from 7 p.m. until 12 p.m., during which time no uploads should occur in order to minimize the burden on the network 52. Beginning at 12 p.m., uploads of event records out of a buffer 42 or 44 would begin. In order to have all STBs 30 upload before 8 a.m., the STBs 30 may be divided into upload groups, e.g., 32, with each group uploading over a selected (e.g., 15 minute) period. To achieve this upload cycle, the following parameters are defined in the FIG. 3 cycle in table XI:

TABLE XI

Upload Cycle Parameters	
Parameter	Definition
Cycle_First_Occurance_Start_Time	12:00 am Jan 1, 1995 + "X" * 15 minutes. "X" staggers each upload group by 15 minutes; X = number of Groups
Cycle Time	24 hours
Upload Duration	15 minutes

A total of four upload cycles will be defined for each group of STBs 30, which allows for weekly uploads or any other combination of cycles to work around peak network 52 load times. 35

STBs 30 can be instructed as to their role in uploading by sending from staging server 70 appropriate commands that are handled by the clickstream upload controller 38. For instance, the following commands may be addressed and sent by staging server 70 to a single or group of STBs 30.

TABLE XII

Clickstream Upload Control Commands	
Octet#	Contents
T 0	Transaction Code MSB = 0 x 80
T 1	Transaction Code LSB = 0 x 10
0	Clickstream Processor Version Number
1	Global Addressable Group Address
	Denotes the group of STBs to field this transaction (b5)
(b1) (b1)	Flag Flag
2	Collection On/Off Key
	Will turn clickstream collection On/Off to a STB or Group of STBs (non-Global only).
3	Perform Upload Now Key
	Will perform an upload on command. Will only upload on command if Global Flag is NOT set.
4	Suppress Upload on Buffer Full
	Will keep the STB or Group from uploading when buffer is full. The STB or Group will only upload on its appointed upload cycle.
5	Upload_Cycle_Definition
	A STB will have 1 to 4 possible upload cycles defined. This will define any one of those cycles.

TABLE XII-continued

Clickstream Upload Control Commands			
Octet#	Contents		
6	Cycle First Occurrence	Year (b8)	Defines the time for first upload in cycle.
	Start Time (Total b48)	Month (b8)	
7	Cycle First Occurrence	Start Time	
	Start Time	Day (b8)	
8	Cycle First Occurrence	Start Time	
	Start Time	Hour (b8)	
9	Cycle First Occurrence	Start Time	
	Start Time	Minute (b8)	
A	Cycle First Occurrence	Start Time	
	Start Time	Second (b8)	
B	Cycle First Occurrence	Start Time	
	Start Time	Hours(0-24) (b8)	Defines a duration of time over which the STB randomizes upload start time.
C	Upload Duration (Total b24)	Minutes(0-59) (b8)	
		Seconds(0-59) (b8)	
D	Upload Duration	Days (0-14) (b8)	Defines the periodicity (mean time) between uploads.
E	Upload Duration	Hours(0-24) (b8)	
F	Cycle Time (Total b32)	Minutes(0-59) (b8)	
		Seconds(0-59) (b8)	
10	Cycle Time		
11	Cycle Time		
12	Cycle Time		

Depending on how the system is configured, the commands instruct STBs 30 to: 1) define the cyclic upload for various groups of or even all STBs 30; 2) require STBs 30 to upload on command/polling control (addressable only); 3) suppress upload when a buffer 42 or 44 fills; or 4) turn on/off event record collection by particular or groups of STBs 30. 35

Event Record Formatting, Upload and Capture

After the upload process triggers, each STB 30 typically initiates upload by first locking the buffer 42 or 44 to be uploaded and then compressing the contents of that buffer 42 or 44. A number of different compression techniques may be used, however, about 50% compression may be achieved with LZW compression utilities. Such compression substantially reduces the load on network 52 caused by numerous STBs 30 uploading event records. Compressed data is divided into transmission "transactions" or "packets" and packet headers are addressed to indicate packet identification, IP destination address, etc. The actual network connection can be initiated by the operating system for the particular STB 30. Persons skilled in the art will recognize that the type of and manner of invoking and implementing the network connection will vary depending upon the type of media delivery network over which system 20 is deployed. 45

For instance, the STB 30 can be configured to insert UDP/IP headers and trailers taken from the RFC 791 or RFC 768 specifications published by the ISO. Each data packet may have UDP/IP protocol built around a Level 1 pass-through header, such as shown in Table XIII below: 55

TABLE XIII

UDP/IP Protocol for Headers				
IP Header				
65	IP Version	Header Length	Type of Service	Total Length

COUNT FOUR: INFRINGEMENT OF U.S. PATENT NO. 6,983,478

27. AT&T incorporates by reference Paragraphs 1 through 26, as if fully set forth herein.

28. United States Patent No. 6,983,478, entitled "Method and System for Tracking Network Use" (hereinafter, the "'478 patent"), duly and legally issued on January 3, 2006 after a full and fair examination. AT&T Intellectual Property I, L.P. is the assignee of all rights, title, and interest in the '478 patent, including the right to sue and recover for all past infringement. A true and correct copy of the '478 patent is attached as Exhibit D.

29. As an example, and not to be limited to only a single infringing product, TiVo has developed, has tested, promotes, markets, uses and sells services for tracking DVR-use patterns, such as “Stop||Watch” and “Power||Watch.” Such services (or the systems used to provide such services) fall within the scope of one or more claims of the ‘478 patent.

30. TiVo has infringed and continues to infringe the '478 patent, by, among other acts, making, using, offering for sale, and/or selling within this Judicial District and elsewhere in the United States, without license or authority by AT&T, services covered by one or more claims of the '478 patent, including, but not limited to, services for tracking DVR-use patterns.

31. As a consequence of TiVo's infringement, AT&T is entitled to recover damages adequate to compensate it for the infringement complained of herein, but in no event less than a reasonable royalty.

32. TiVo has caused and will continue to cause AT&T substantial damage and irreparable injury by virtue of its past and continuing infringement of the '478 patent. TiVo will continue to infringe the '478 patent unless enjoined by this Court. AT&T will suffer further damage and irreparable injury unless and until TiVo is enjoined by this Court from continuing such infringement.

JURY DEMAND

AT&T hereby demands a jury trial on all issues so triable.

TABLE XIII-continued

UDP/IP Protocol for Headers			
Identification	Flags	Fragment Offset	
Time to Live	Protocol	Header Checksum	
Source IP Address			
Destination IP Address			
UDP Header			
Source Port		Destination Port	
Length		Checksum	

In the embodiment shown in the Figures, the clickstream processor 34 will identify a particular VSP—Video Service Provider, which is an entity connecting to network 52 to distribute services—like VSP 66 shown in FIG. 4B, as the destination of these data packets. All of the data to be uploaded appears as “payload” to the STB 30, the signaling network 52, the network management controller 50, and the event capture process 71 on the staging server 70. After an appropriate header and trailer inserted at the STB 30, the upload data packet may have the format shown in Table XIV:

TABLE XIV

Clickstream Upload Data Packet	
Octet#	Contents
T 0	Transaction Code MSB = 0x80
T 1	Transaction Code LSB = 0x18
0	Clickstream Processor Version Number
1	Upload Sequence Number
0x02	Clickstream Upload Buffer Data Structure (as shown in Table I and X). The data may be broken up into as many reverse path transactions as necessary to complete data upload.
0xFA	

Providing two buffers 42, 44 allows event record collection to continue during upload. Assuming buffer 42 is being uploaded, if the second buffer 44 fills during the upload process, a buffer overrun condition occurs. To account for such an occurrence, the buffer trailer record sent during upload from STBs 30 may denote such an error condition. The structure of the buffer trailer record may take the form as shown in Table XV below and include a time stamp, assigned application identification, length and upload code.

TABLE XV

Buffer Trailer Record	
	Size
Timestamp	6 bytes
Assigned Application ID	16 bits
Number Bytes to Follow (length)	8 bits
Upload Status Code	8 bits

These upload status codes identify the stage of the upload process at the time a buffer 42 or 44 overflow occurred. Thus, some possible upload codes could include: upload not used, upload in progress, upload completed but no acknowledgment received, upload completed but only partial acknowledgment received or no upload attempted. This will let the staging server 70 know that STB 30 event records are missing beginning at that time. Also, receiving a buffer overrun record informs the staging server 70 that buffer 42 or 44 sizes have not been set appropriately. Buffer 42 or 44

sizes can then be reset and released to the system 20 as an update or released to a particular STB 30 by sending it an appropriate command.

Note that the packetization description above is for one embodiment of the system 20. However, generally, to upload collected event records, STBs 30 can initiate whatever “upstream” data transmission process used by the interactive, cable television or other media delivery system with which the system 20 is used. That process will upload the event records in the appropriate system format.

In any event, for system 20, clickstream data packets are uploaded to the staging server 70 over a slotted-ALOHA (a contention-based standard transport protocol) data transmitter of the STB 30. Data acknowledgments from the staging server 70 are sent; each is addressed to particular STBs 30. The frequency and period of data acknowledgments may be determined by considering network error rates, network packet error rates and causes of those types transmission errors.

FIGS. 4A and 4B show in greater detail the clickstream data flow through the system 20. Briefly, FIG. 4A shows that clickstream packets of event records are transmitted from each STB 30 to the network management controller 50, which acts as a video service provider router. From the network management controller 50, which manages traffic over network 52, packets are forwarded via the network 52, video server 60 and video server control 56 to the staging server 70, which couples to MKIS 100 and analysis engine 110. Thus, event records collected and buffered at STBs 30 are transmitted to the staging server 70 for collection and analysis.

FIG. 4B shows this process in more detail and also describes an event records capture process 71 at staging server 70.

As noted, once a buffer 42 or 44 fills or the clickstream processor 34 decides to upload data for other reasons (time expiration, low system utilization, commanded upload, etc.), the buffer 42 or 44 will be formatted, compressed and then uploaded through the system 20 to the staging server 70. The upstream data packets may travel from the network management controller 50 across the distribution network 52 to video server 60 through a process called IP (“Internet Protocol”) tunneling, which is essentially automatic IP routing based upon information in the packet payload. The same process can be used to route packets through network 52 directly to staging server 70 without going through video server 60. FIG. 4B shows that, at video server 60, an L1 pass-through process 63 uses a VSP routing table 67 to associate destination IP addresses with corresponding tags inserted in the received data packets. This process re-directs the data packets to the application server 66 L1 pass-through process 63 by associating the tags with the appropriate listed destination—here, the application server 66. The L1 pass-through process 63 on application server 66 performs a similar function with the data packets, routing them based on a payload identifier (transaction code or other) to an event record capture (“ECAP”) open server process 71 on the staging server 70.

When the ECAP process 71 receives a clickstream data packet, it accepts the data packet and correlates the source address of the data packet with an upload session already in progress with a particular STB 30. If there is currently no upload in progress with that STB 30, then one is considered to be initiated. ECAP process 71 processes the upload of data in accordance with the particular protocol needed for the system 20. After receipt of all clickstream data packets associated with the upload from a particular STB 30, the

ECAP process 71 sequences the packets into proper order (particular packets may have arrived out of their original transmission sequence because of transmission delays in network 52), decompresses the packets, eliminates transport overhead (e.g., trailers, headers, etc.) and stores them, such as in a flat file, for later analysis. At the end of a selected period, like 24 hours, the file is closed and a new one is opened, which allows a subsequent merge and parse process to batch process discrete files that cover discrete time periods. Immediately after initiation of and during the ECAP process 71, an operation log is opened to record information about the initiation and termination of each upload session and any errors.

As shown in FIG. 5, staging server 70 will formulate and send a data acknowledgment to each STB 30 engaged in the upload process. One method of doing so is to send acknowledgments as addressable downstream level one pass-through transactions over network 52 to the STB 30. Such data acknowledgments provide redundant error correction because failure to receive them may alert STB 30 to a possible transmission error.

Merging and Parsing

FIGS. 6A and 6B show an overview of the merging and parsing process and FIG. 7 shows sample results following that process. Briefly, the aim of the merge and parse process is to merge each STB 30's event records with various "metadata." "Metadata" refers to (1) programming of virtually any type shown on system 20 including the time and broadcast or cable network providing such programming or (2) interactive applications invoked by subscribers. For instance, metadata includes the following sources of data: EPG broadcast programming schedule data 82, broadcast advertising schedule data 84, local advertising schedule data or session-services advertising schedule data 86 and session-services programming schedule data 88. (Session-services advertising refers to advertising inserted by video server 60 during particular interactive sessions with the subscriber that are the session-services programming).

Collectively, all of this data enters into a merge and parse engine 90 that creates an event timeline 92 for each STB 30. Merge and parse engine 90 may be deployed upon staging server 70 or the MKIS system 100. So deploying merge and parse engine 90 on staging server 70 allows collected event records to be merged and parsed. The resulting event timelines 92 can be sent to MKIS system 100 for further analysis.

Timeline 92 provides a snapshot of activity on a particular STB 30 for a selected period (e.g., 24 hours) or for a selected event—for instance, a timeline 92 would be created for each STB 30 tuning to a particular show or shows (e.g., a pay per view fight) that may occur over a selected period. Timeline 92 is created by merging event records with metadata about programming available over the network for the selected time period.

To merge that data, proper priority must be assigned to data that otherwise may be conflicting. For instance, broadcast advertising data 84 may indicate that a certain national ad was run at Time A. On the other hand, if the system 20 is an interactive system and the interactive server provided a targeted advertisement ("ad") also at Time A, as indicated by session-services advertising data 86, that targeted ad was inserted over the national ad at Time A. Thus, by assigning session-services advertising data 86 a priority higher than national broadcast advertising data 84, the merge and parse engine 90 is able to create an accurate timeline 92 of programming delivered to a particular STB 30. Similarly,

even a traditional cable or wireless cable network requires priority assignments. Typically, local cable operators typically are allowed to insert local ads over certain national ads (assuming they can sell that local ad time).

FIG. 6B depicts such priority assignments. FIG. 6B shows several sources of data, such as EPG metadata, National and Local Insert ad metadata and Interactive Sessions metadata. EPG metadata is usually very broad—for instance, showing a football game on channel 1 from 1:00 to 4:00 p.m. Thus, EPG metadata is assigned a priority lower than that of national ad metadata because a particular national ad will be overlayed into a particular time slot broadly defined by the EPG. In turn, local insert ad metadata trumps national ad metadata because the national ad metadata may not account for situations where a local network or affiliate inserts a local ad over the national ad scheduled for a particular timeslot. Finally, interactive sessions metadata, which reflects subscriber selections, has the highest priority as it shows the subscriber stopped watching a particular channel and instead invoked an interactive session.

Applying these priority rules produces a timeline 94 for each subscriber. Additional filtering criteria 94 are applied by the merge and parse engine 90 in order to generate a further refined timeline 94, as depicted in FIG. 6A. For example, event records may include such highly granular and specific information as the number of volume ups or channel ups that a particular subscriber entered. One set of filtering criteria 94 may ensure that the timeline 92 includes only channels that were viewed for more than a threshold (e.g., 15 seconds) time period. This eliminates any very fast channel changes made by the subscribers, thereby simplifying the event timeline 92 because event records that do not meet the criteria 94 are filtered out of the event timeline 92.

Merge and parse engine 90 also may apply other criteria to the filtered timeline 94 (or the original timeline 92), as shown in FIG. 6. Specifically, advertisers may wish to apply "view" and "watch" criteria 96. This criteria 96 will identify those programs and advertisements that are "viewed" by subscribers for less than a certain threshold amount of time. Programming seen by subscribers for more than that threshold, would be identified as "watched" programming. For example, for a 30 second ad, the threshold might be 15 seconds. If a subscriber was tuned to a channel displaying that ad for less than 15 seconds he would be deemed to have simply "viewed" that ad; on the other hand, if the subscriber was tuned to the channel carrying that ad for 25 seconds of the ad's length, he would be deemed to have "watched" it. This criteria 96 allows system 20 operators to charge more for "watched" ads versus those that are merely "viewed." Similar criteria can be applied against programming in order to more accurately gauge ratings. Thus, for a 30 minute program, if a user was tuned to that program for less than 10 minutes, the view and watch criteria 96 may decide that the program was only "viewed." In any event, applying the view and watch criteria 96, merge and parse engine 90 creates "view" and "watch" lists 98 that are useful for the system 20 operator and advertisers who contract with system 20 operator.

Note also that other criteria than simply how much time a tuned to a particular channel may be included in the view and watch criteria 96. For instance another criteria may be volume level. If a viewer was tuned to a channel for the full thirty second length of an ad but hit the mute button or

changed the volume below a certain threshold for that ad, view and watch criteria 96 may classify that ad as a "viewed" ad.

Generally, merging and parsing should be done on discrete segments of data, such as 24 hour segments, as soon as possible in order to minimize the occurrence of un-resolved events. In other words, discrete events are simply pieces of the entire picture. To analyze only several hours of clickstream event data would not allow determination of such things as programming "watched" versus "viewed."

FIG. 7 shows a sample merge of event records or click-stream data 80, EPG data 82 from Prevue or a similar service and broadcast advertising data 84 that creates a clickstream timeline 92, which shows both the channels selected by a subscriber and the content displayed on those channels while the subscriber watched them.

A timeline 94 for each STB 30 is built and uploaded by staging server 70 to the MKIS database 100 or a third party analysis engine and database 110, either of which may store demographics and be used to run queries against the event timelines 94 and those demographics. Combining the timelines 94 with demographics information allows for even more detailed and granular information about subscribers and their viewing habits. For instance, consider the following examples:

EXAMPLE 1

Widget Co. has ten different advertisements that it has been running on system 20. Widget Co. wishes to know whether subscribers are "viewing" or "watching" particular ads. Because of the detailed information captured by the system 20 of the present invention, a query can be formulated to determine (a) which subscribers "watched" particular 30 second advertisements for greater than 15 seconds versus (b) which subscribers simply "viewed" the ad, for less than 15 seconds.

EXAMPLE 2

When event timelines 94 (or view and watch lists 98) are loaded into MKIS 100 or analysis engine 110, the same query can be run for a particular demographic group. For instance, Widget Co. wishes to know which particular ads its primary customer base, baby boomers between ages 40 and 50 and with income over \$50,000 per year, "watched" versus "viewed" their advertisements.

Obviously, the system 20 can also be modified to target ads to particular demographic households based on feedback from parsed and merged data. Then, event records occurring after those targeted ads are broadcast over system 20 can be checked to determine whether the particular demographic market targeted watched or viewed the advertisement.

The foregoing is provided to explain and disclose preferred embodiments of the present invention, modifications to which may be made that still fall within the following claims. For instance, the architecture and programming of the system may be modified. Or, a variety of different manufacturers' servers, set top boxes or databases may be configured in order to implement the system. The particular identification codes and allocated sizes show in the tables and described herein may also be greatly modified. Further modifications and adaptations to the described embodiments will be apparent to those skilled in the art and may be made without departing from the scope or spirit of the invention and the following claims.

What is claimed is:

1. A system for collecting and processing information about subscribers' selection and use of programming distributed over a media delivery network, the system comprising:
 - a) merge processor coupled via means for communication to
 - b) a plurality of set top boxes, each comprising a processor for
 - (1) collecting a plurality of event records that describe selected commands from a subscriber to a particular set top box and
 - (2) transmitting event records to the merge processor;
 - c) wherein the merge processor forms an event timeline describing a subscriber's selection of distributed programming for a discrete time period by merging the event records with programming data describing programming available via the media delivery system.
2. A system according to claim 1 wherein the programming data comprises data collected from at least two sources selected from the group consisting of: a broadcasting schedule source, a national advertising schedule source, a local advertising schedule source and an interactive application use schedule source.
3. A system according to claim 1 wherein each set top box further comprises a plurality of applications capable of being invoked by a subscriber.
4. A system according to claim 3 wherein each event record comprises:
 - (1) an application identifier corresponding to the application associated with the recorded event;
 - (2) an event identification code; and
 - (3) a time stamp associated with the initiation of the event.
5. A system according to claim 4 wherein each application creates an event upon detection of selected commands from the subscriber.
6. A system according to claim 1 further comprising a buffer for storing the event records before transmission.
7. A system according to claim 1 wherein the merge processor forms an event timeline for each of the plurality of set top boxes.
8. A system according to claim 7 further comprising an analysis engine for correlating the event timelines with demographics information describing the subscribers.
9. A method for journaling information about subscriber use of a media delivery network for delivering programming and a merge processor for analyzing the resulting journaled information, the method comprising the steps of:
 - a) collecting information about a plurality of subscribers' use of a media delivery network, the collecting step comprising:
 - i) identifying commands of interest from each subscriber;
 - ii) forming event records that record at least the commands of interest and a time associated with the command;
 - b) transmitting event records to the merge processor;
 - c) merging the event records with data describing the programming delivered over the media network in order to form event timelines of which describes the programming selected by a particular subscriber over a discrete time period.

10. A method according to claim 9 wherein the identifying step comprises the step of correlating each command of interest with a global table comprising identification codes.

11. A method according to claim 9 further comprising the step of filtering the event timelines in order to classify subscribers' viewing patterns into at least two categories.

12. A method according to claim 11 wherein the first category comprises programming watched by a subscriber for greater than a selected threshold percent of the total program length.

13. A system for determining the viewing habits of subscribers to a media delivery network for delivering programming, the system comprising:

- a) a collector for collecting event records describing subscribers' selection and use of programming;
- b) means, coupled to the collector, for communicating event records to
- c) a merge processor for processing the event records to form a selected subscriber an event timeline describing the programming delivered to a selected subscriber over a particular time period via the media delivery network;
- d) means for storing demographics information about selected groups of subscribers; and

5 e) wherein the merge processor forms a plurality of event timelines and correlates the demographics information with the event timelines.

14. A system according to claim 13 in which the merge processor applies filtering criteria to the event records to determine the programming watched by a subscriber for greater than a selected percent of the programming.

10 15. A system according to claim 14 in which the collector is deployed upon a set top box that is associated with a display device for displaying delivered programming.

15 16. A system according to claim 15 in which the subscriber controls the set top box via a remote device in order to invoke and run a variety of applications and the collector forms event records by:

- a) identifying a code that corresponds to a command of interest entered by a selected subscriber; and
- b) storing in a buffer, associated with the collector, an event record comprising
 - (1) the code corresponding to the command; and
 - (2) a time stamp.

* * * * *

PRAYER

WHEREFORE, AT&T prays for relief against TiVo as follows:

a. A judgment that TiVo has infringed the claims of U.S. Patent Nos. 5,809,492; 5,922,045; 6,118,976; and 6,983,478 (collectively, “the patents-in-suit”);

b. A permanent injunction enjoining TiVo, its officers, agents, servants, employees, representatives, licensees, successors, assigns, and those persons in active concert or participation with any of them, from making, using, selling, offering to sell, and importing products, services, or processes that infringe the patents-in-suit;

c. Awarding AT&T damages adequate to compensate for the infringement by TiVo, but in no event less than a reasonable royalty for the use made of the invention by TiVo, together with interest and costs under 35 U.S.C. § 284;

- d. Awarding pre-judgment and post-judgment interest on the damages assessed; and
- e. Awarding to AT&T such other and further relief as the Court deems just.

Dated: March 12, 2010

Respectfully submitted,

~~BAKER BOTTS L.L.P.~~

Bryant C. Boren Jr.
Attorneys for Plaintiffs
AT&T Intellectual Property I, L.P. and
AT&T Intellectual Property II, L.P.